arXiv:1703.01398v2 [cs.RO] 7 Mar 2017

Sparse Depth Sensing for Resource-Constrained Robots

Fangchang Ma, Luca Carlone, Ulas Ayaz, Sertac Karaman

***Abstract*— We consider the case in which a robot has to navigate in an unknown environment, but does not have enough on-board power or payload to carry a traditional depth sensor (e.g., a 3D lidar) and thus can only acquire a few (point-wise) depth measurements. Therefore, we address the following question: *is it possible to reconstruct the geometry of an unknown environment using sparse and incomplete depth measurements?* Reconstruction from incomplete data is not possible in general, but when the robot operates in man-made environments, the depth exhibits some regularity (e.g., many planar surfaces with few edges); we leverage this regularity to infer depth from a small number of measurements. Our first contribution is a formulation of the depth reconstruction problem that bridges robot perception with the *compressive sensing* literature in signal processing. The second contribution includes a set of formal results that ascertain the exactness and stability of the depth reconstruction in 2D and 3D problems, and completely characterize the geometry of the signals that we can reconstruct. Our third contribution is a set of practical algorithms for depth reconstruction: our formulation directly translates into algorithms for depth estimation based on convex programming. In real-world problems, these convex programs are very large and general-purpose solvers are relatively slow. For this reason, we discuss ad-hoc solvers that enable fast depth reconstruction in real problems. The last contribution is an extensive experimental evaluation in 2D and 3D problems, involving Monte Carlo runs on simulated instances and testing on multiple real datasets. Empirical results confirm that the proposed approach ensures accurate depth reconstruction, outperforms interpolation-based strategies, and performs well even when the assumption of structured environment is violated.**

Supplemental Material

- Video demonstrations: `https://youtu.be/vE56akCGeJQ`
- Source code for the algorithms: `https://github.com/sparse-depth-sensing/sparse-depth-sensing.git`

I. Introduction

Recent years have witnessed a growing interest towards miniaturized robots (e.g., the *RoboBee* [75], *Piccolissimo* [1], the *DelFly* [17], [16], the *Black Hornet Nano* [2], *Salto* [30]). These robots are usually palm-sized (or even smaller), can be deployed in large numbers, and provide a new perspective on societally relevant applications, including artificial pollination, environmental monitoring, and disaster response. Despite the rapid development and recent success in control, actuation, and manufacturing of miniature robots, on-board sensing and perception capabilities for such robots remain a relatively unexplored, challenging open problem. These small platforms have extremely limited payload, power, and on-board computational resources, and thus preventing the use of standard sensing and computation paradigms.

F. Ma, L. Carlone, U. Ayaz, and S. Karaman are with the Laboratory for Information & Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA, {fcma,lcarlone,uayaz,sertac}@mit.edu

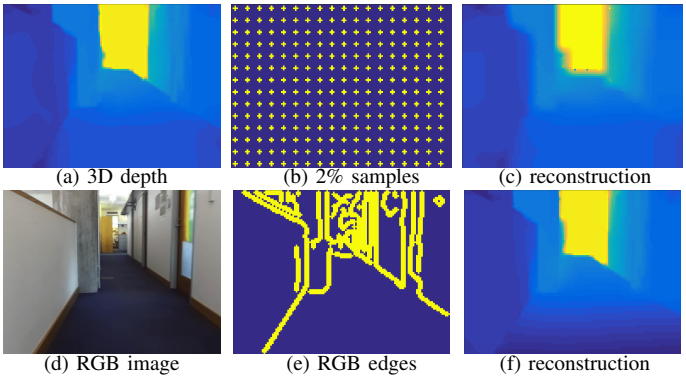

Fig. 1. We show how to reconstruct an unknown depth profile (a) from a handful of samples (b). Our reconstruction is shown in (c). Our results also apply to traditional stereo vision and enable accurate reconstruction (f) from few depth measurements (e) corresponding to the edges in the RBG image (d). Figures (a) and (d) are obtained from a ZED stereo camera.

In this paper we explore novel sensing techniques for miniaturized robots that cannot carry standard sensors. In the last two decades, a large body of robotics research focused on the development of techniques to perform inference from data produced by "information-rich" sensors (e.g., high-resolution cameras, 2D and 3D laser scanners). For instance, a variety of approaches has been proposed to perform geometric reconstruction using these sensors (see [50], [44], [74] and the references therein). On the other extreme of the sensor spectrum, applications and theories have been developed to cope with the case of minimalistic sensing [66], [18], [70], [68]. In this latter case, the sensor data is usually not metric (i.e., the sensor cannot measure distances or angles) but instead binary in nature (e.g., binary detection of landmarks), and the goal is to infer only the topology of the (usually planar) environment rather than its geometry. This work studies a relatively unexplored region between these two extremes of the sensor spectrum.

Our goal is to design algorithms (and lay the theoretical foundations) to reconstruct a depth profile (i.e., a laser scan in 2D, or a depth image in 3D, see Fig. 1) from sparse and incomplete depth measurements. Contrary to the literature on minimalistic sensing, we provide tools to recover complete geometric information, while requiring much fewer data points compared to standard information-rich sensors. This effort complements recent work on hardware and sensor design, including the development of lightweight, small-sized depth sensors. For instance, a number of ultra-tiny laser range sensors are being developed as research prototypes (e.g., the dime-sized, 20-gram laser [13], and an even smaller

lidar-on-a-chip system with no moving parts [57]), while some other distance sensors have already been released to the market (e.g., the TeraRanger's single-beam, 8-gram distance sensor [67], and the LeddarVu's 8-beam, 100-gram laser scanner [36]). These sensors provide potential hardware solutions for sensing on micro (or even nano) robots, as well as other compact-sized mobile platforms. Although these sensors meet the requirements of payload and power consumption of miniature robots, they only provide very sparse and incomplete depth data, in the sense that the raw depth measurements are extremely low-resolution (or even only a few beams). In other words, the output of these sensors cannot be utilized directly in high-level tasks (e.g., object recognition and mapping), and the need to reconstruct a complete depth profile from such sparse data arises.

Contribution. We address the following question: *is it possible to reconstruct a complete depth profile from sparse and incomplete depth samples?* In general, the answer is negative, since the environment can be very adversarial (e.g., 2D laser scan where each beam is drawn randomly from a uniform distribution), and it is impossible to recover the depth from small set of measurements. However, when the robot operates in structured environments (e.g., indoor, urban scenarios) the depth data exhibits some regularity. For instance, man-made environments are characterized by the presence of many planar surfaces and a few edges and corners. This work shows how to leverage this regularity to recover a depth profile from a handful of sensor measurements. Our overarching goal is two-fold: to establish theoretical conditions under which depth reconstruction from sparse and incomplete measurements is possible, and to develop practical inference algorithms for depth estimation.

Our first contribution, presented in Section IV, is a general formulation of the depth estimation problem. Here we recognize that the "regularity" of a depth profile is captured by a specific function (the ℓ_0-norm of the 2^{nd}-order differences of the depth profile). We also show that by relaxing the ℓ_0-norm to the (convex) ℓ_1-norm, our problem falls within the cosparsity model in *compressive sensing* (CS). We review related work and give preliminaries on CS in Sections II-III.

The second contribution, presented in Section V, is the derivation of theoretical conditions for depth recovery. In particular, we discuss the conditions under which reconstruction of a profile from incomplete measurements is possible, investigate the robustness of depth reconstruction to noise, and provide bounds on the reconstruction error. Contrary to the existing literature in CS, our conditions are geometric (rather then algebraic) and provide readily usable information to guide sampling strategy.

Our third contribution, presented in Section VI, is algorithmic. We discuss practical algorithms for depth reconstruction, including different variants of the proposed optimization-based formulation, and solvers that enable fast depth recovery. In particular, we discuss the application of a state-of-the-art solver for non-smooth convex programming, called NESTA [4], to our depth estimation problem.

Our fourth contribution, presented in Section VII, is an extensive experimental evaluation, including Monte Carlo runs on simulated data and testing with real sensors. The experiments confirm our theoretical findings and show that our depth reconstruction approach is extremely resilient to noise and works well even when the regularity assumptions are violated. We discuss many applications for the proposed approach. Besides our motivating scenario of navigation with miniaturized robots, our approach finds application in several endeavors, including multi-frame depth reconstruction, data compression, and super-resolution depth estimation.

Section VIII draws conclusions and discusses future research. All proofs, together with extra visualizations, are given in the appendix of the paper.

II. Related Work

This work intersects several lines of research across fields.

Minimalistic Sensing. Our study of depth reconstruction from sparse sensor data is related to the literature on *minimalistic sensing*. Early work on minimalistic sensing includes contributions on sensor-less manipulation [27], robot sensor design [7], [51], and target tracking [64]. Suri *et al.* [66], Derenick *et al.* [18], and Tovar *et al.* [3] use binary measurements of the presence of landmarks to infer the topology of the environment. Marinakis *et al.* [42], [41] reconstruct the topology of a sensor network from unlabeled observations from a mobile robot. O'Kane *et al.* [52] and Erickson *et al.* [22] investigate a localization problem using contact sensors. Tovar *et al.* [71] use depth discontinuities measurements to support exploration and search in unknown environments. Tovar *et al.* [69], [68] propose a combinatorial filter to estimate the path (up to homotopy class) of a robot from binary detections. Milford [43] addresses minimality of information for vision-based place recognition.

Sensing and perception on miniaturized robots. A fairly recent body of work in robotics focuses on miniaturized robots that draw inspiration from small animals and insects. Most of the existing literature focuses on the control of such robots, either open-loop or based on information from external sensors. However, there has been relatively little work on onboard sensing and perception. For example, the *Black Hornet Nano* [2] is a military-grade micro aerial vehicle equipped with three cameras but with basically no autonomy. *Salto* [30] is Berkeley's 100g legged robot with agile jumping skills. The jump behavior is open-loop due to lack of sensing capabilities, and the motion is controlled by a remote laptop. The *RoboBee* [75] is an 80-milligram, insect-scale robot capable of hovering motions. The state estimation relies completely on an external array of cameras. *Piccolissimo* [1] is a tiny, self-powered drone with only two moving parts, controlled completely by an external, hand-held infrared signal. The *DelFly Explorer* [17], [16] is a 20-gram flying robot with an onboard stereo vision system. It is capable of producing a coarse depth image at 11Hz and is thus one of the first examples of miniaturized flying robot with basic obstacle avoidance capabilities.

Fast Perception and Dense 3D Reconstruction. The idea of leveraging priors on the structure of the environ-

ment to improve or enable geometry estimation has been investigated in early work in computer vision for single-view 3D reconstruction and feature matching [34], [31]. Early work by Faugeras *et al.* [24] addresses *Structure from Motion* by assuming the environment to be piecewise planar. More recently, Pillai *et al.* [55] propose an approach to speed-up stereo reconstruction by computing the disparity at a small set of pixels and considering the environment to be piecewise planar elsewhere. Dame *et al.* [15] combine live dense reconstruction with shape-priors-based 3D tracking and reconstruction. Estellers *et al.* [23] propose a regularization based on the structure tensor to better capture the local geometry of images. Lu *et al.* [39] produce high-resolution depth maps from subsampled depth measurements by using segmentation based on both RGB images and depth samples. Piniés *et al.* [56] compute a dense depth map from a sparse point cloud. This work is related to our proposal with three main differences. First, the work [56] uses an energy minimization approach that requires parameter tuning (the authors use Bayesian optimization to learn such parameter); our approach is parameter free and only assumes bounded noise. Second, we use a 2^{nd}-order difference operator to promote depth regularity, while [56] considers alternative costs, including nonconvex regularizers. Finally, by recognizing connections with the *cosparsity* model in CS, we provide theoretical foundations for the reconstruction problem.

Map Compression. Our approach is also motivated by the recent interest in *map compression*. Nelson and Michael [47] propose a compression method for occupancy grid maps, based on the *information bottleneck* theory. Ramos *et al.* [58], [53] use Gaussian processes to improve 2D mapping quality from smaller amount of laser data. Im *et al.* [32] investigate wavelet-based compression techniques for 3D point clouds. Ruhnke *et al.* [62], [63] discuss point cloud compression techniques based on sparse coding.

Compressive Sensing (CS). Finally, our work is related to the literature on *compressive sensing* [10], [19], [11], [25]. While Shannon's theorem states that to reconstruct a signal (e.g., a depth profile) we need a sampling rate (e.g., the spatial resolution of our sensor) which must be at least twice the maximum frequency of the signal, CS revolutionized signal processing by showing that a signal can be reconstructed from a much smaller set of samples if it is *sparse* in some domain. CS mainly invokes two principles. First, by inserting *randomness* in the data acquisition, one can improve reconstruction. Second, one can use ℓ_1-minimization to encourage sparsity of the reconstructed signal. Since its emergence, CS impacted many research areas, including image processing (e.g., inpainting [37], total variation minimization [46]), data compression and 3D reconstruction [72], [60], [20], inverse problems and regularization [73], matrix completion [59], and single-pixel imaging techniques [54], [21], [65]. While most CS literature assumes that the original signal z is sparse in an appropriate domain, i.e., $z = Dx$ for some a matrix D and a sparse vector x (this setup is usually called the *synthesis model*), very recent work considers the case in which the signal becomes sparse after a transformation is applied (i.e., given a matrix D, the vector Dz is sparse). The latter setup is called the *analysis (or cosparsity) model* [45], [33]. An important application of the analysis model is total variation minimization which is ubiquitous in image processing [46], [61]. In a hindsight we generalize total variation (which applies to piecewise constant signals) to deal with piecewise *linear* functions.

A closely related work in signal processing is the paper by Liu *et al.* [38], which uses a combined dictionary of wavelets and contourlets for sparse representation of disparity maps and propose a new ADMM algorithm for reconstruction from sparse samples. In comparison, we use a much more intuitive geometric model to represent the sparsity of the signal and provide stronger theoretical guarantees. Our formulation enables the use of fast solvers for depth reconstruction. Finally, our approach is able to operate with much fewer samples compared to [38].

III. Preliminaries and Notation

We use uppercase letters for matrices, e.g., $D \in \mathbb{R}^{p\times n}$, and lowercase letters for vectors and scalars, e.g, $z \in \mathbb{R}^n$ and $a \in \mathbb{R}$. Sets are denoted with calligraphic fonts, e.g., $\mathcal{M}$. The cardinality of a set $\mathcal{M}$ is denoted with $|\mathcal{M}|$. For a set $\mathcal{M}$, the symbol $\overline{\mathcal{M}}$ denotes its complement. For a vector $z \in \mathbb{R}^n$ and a set of indices $\mathcal{M}$, $z_{\mathcal{M}}$ is the subvector of z corresponding to the entries of z with indices in $\mathcal{M}$. In particular, z_i is the i-th entry. The symbols $\mathbf{1}$ (resp. $\mathbf{0}$) denote a vector of all ones (resp. zeros). The support set of a vector is denoted with

$$\operatorname{supp}(z) = \{i \in \{1,\dots,n\} : z_i \neq 0\}.$$

We denote with $\|z\|_2$ the Euclidean norm and we also use the following norms:

$$\|z\|_\infty \doteq \max_{i=1,\dots,n} |z_i| \quad (\ell_\infty \text{ norm}) \tag{1}$$

$$\|z\|_1 \doteq \sum_{i=1,\dots,n} |z_i| \quad (\ell_1 \text{ norm}) \tag{2}$$

$$\|z\|_0 \doteq |\operatorname{supp}(z)| \quad (\ell_0 \text{ pseudo-norm}) \tag{3}$$

Note that $\|z\|_0$ is simply the number of nonzero elements in z. The sign vector $\operatorname{sign}(z)$ of $z \in \mathbb{R}^n$ is a vector with entries:

$$\operatorname{sign}(z)_i \doteq \begin{cases} +1 & \text{if } z_i > 0 \\ 0 & \text{if } z_i = 0 \\ -1 & \text{if } z_i < 0 \end{cases}$$

For a matrix D and an index set $\mathcal{M}$, we denote with $D_{\mathcal{M}}$ the submatrix of D containing only the *rows* of D with indices in $\mathcal{M}$. Sometimes, given two sets $\mathcal{I}$ and $\mathcal{J}$, we also use the notation $D_{\mathcal{I},\mathcal{J}}$ which is a sub-matrix of D including only rows in $\mathcal{I}$ and columns in $\mathcal{J}$. The identity matrix is denoted with $\mathbf{I}$. Given a matrix $D \in \mathbb{R}^{p\times n}$, we use the following matrix operator norm $\|D\|_{\infty\to\infty} \doteq \max_{i=1,\dots,p} \|D_i\|_1$.

In the rest of the paper we use the *cosparsity model* where we assume that the application of an *analysis operator* D produces a vector with small number of nonzero entries. The following definitions formalize this concept.

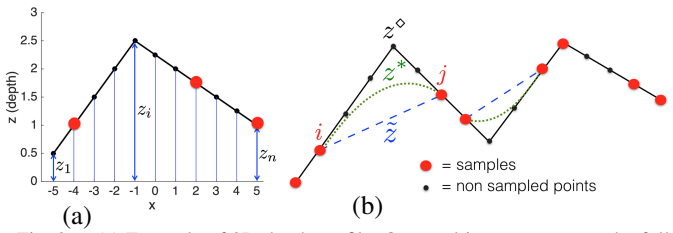

Fig. 2. (a) Example of 2D depth profile. Our goal is to reconstruct the full signal (black solid line) from sparse samples (red dots); (b) When we do not sample the corners and the neighboring points, ($\mathrm{L1}_D$) admits multiple minimizers, which obey the conditions of Proposition 14.

Definition 1 (Cosparsity). *A vector $z \in \mathbb{R}^n$ is said to be* cosparse *with respect to a matrix $D \in \mathbb{R}^{p\times n}$ if $\|Dz\|_0 \ll p$.*

Definition 2 (D-support and D-cosupport). *Given a vector $z \in \mathbb{R}^n$ and a matrix $D \in \mathbb{R}^{p\times n}$, the* D-support *of z is the set of indices corresponding to the nonzero entries of Dz, i.e., $\mathcal{I} = \mathrm{supp}(Dz)$. The* D-cosupport *$\mathcal{J}$ is the complement of $\mathcal{I}$, i.e., the indices of the zero entries of Dz.*

IV. Problem Formulation: Sparse Depth Sensing and Reconstruction

Our goal is to reconstruct 2D depth profiles (i.e., a scan from a 2D laser range finder) and 3D depth profiles (e.g., a depth image produced by a kinect or a stereo camera) from partial and incomplete depth measurements. In this section we formalize the depth reconstruction problem, by first considering the 2D and the 3D cases separately, and then reconciling them under a unified framework.

A. 2D Depth Reconstruction

In this section we discuss how to recover a 2D depth profile $z^\diamond \in \mathbb{R}^n$. One can imagine that the vector $z^\diamond$ includes (unknown) depth measurements at discrete angles; this is what a standard planar range finder would measure.

In our problem, due to sensing constraints, we do not have direct access to $z^\diamond$, and we only observe a subset of its entries. In particular, we measure

$$y = Az^\diamond + \eta \qquad \text{(sparse measurements)} \tag{4}$$

where the matrix $A \in \mathbb{R}^{m\times n}$ with $m \ll n$ is the *measurement matrix*, and η represents measurement noise. The structure of A is formalized in the following definition.

Definition 3 (Sample set and sparse sampling matrix). *A sample set $\mathcal{M} \subseteq \{1,\dots,n\}$ is the set of entries of the profile that are measured. A matrix $A \in \mathbb{R}^{m\times n}$ is called a* (sparse) sampling matrix *(with sample set $\mathcal{M}$), if $A = \mathbf{I}_{\mathcal{M}}$.*

Recalling that $\mathbf{I}_{\mathcal{M}}$ is the subset of rows of the identity matrix at indices $\mathcal{M}$, it follows that $Az = z_{\mathcal{M}}$, i.e., the matrix A selects a subset of entries from z. Since $m \ll n$, we have much fewer measurements than unknowns, hence $z^\diamond$ cannot be recovered from y, without further assumptions.

In this paper we assume that the profile is $z^\diamond$ sufficiently regular, meaning that it contains only a few "corners", e.g., Fig. 2(a). Corners are produced by changes of slope: considering 3 consecutive points at coordinates (x_{i-1}, z_{i-1}), (x_i, z_i), and (x_{i+1}, z_{i+1}),[1] there is a corner at i if $\frac{z_{i+1}-z_i}{x_{i+1}-x_i} - \frac{z_i-z_{i-1}}{x_i-x_{i-1}} \neq 0$. In the following we assume that $x_i - x_{i-1} = 1$ for all i: this comes without loss of generality since the full profile is unknown and we can reconstruct it at arbitrary resolution (i.e., at arbitrary x). We can then formalize the definition of "corner" as follows.

Definition 4 (Corner set). *Given a 2D depth profile $z \in \mathbb{R}^n$, the* corner set *$\mathcal{C} \subseteq \{2,\dots,n-1\}$ is the set of indices i such that $z_{i-1} - 2z_i + z_{i+1} \neq 0$.*

Intuitively, $z_{i-1} - 2z_i + z_{i+1}$ is the discrete equivalent of the 2nd-order derivative at z_i. We call $z_{i-1} - 2z_i + z_{i+1}$ the *curvature* at sample i: if this quantity is zero, the neighborhood of i is flat (the three points are collinear); if it is negative, the curve is locally concave; if it is positive, it is locally convex. To make notation more compact, we introduce the *2nd-order difference operator*:

$$D \doteq \begin{bmatrix} 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & 1 & -2 & 1 & \dots & 0 \\ \vdots & 0 & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & -2 & 1 \end{bmatrix} \in \mathbb{R}^{(n-2)\times n} \tag{5}$$

Then a profile with only a few corners is one where $Dz^\diamond$ is sparse. In fact, the ℓ_0-norm of $Dz^\diamond$ counts exactly the number of corners of a profile:

$$\|Dz^\diamond\|_0 = |\mathcal{C}| \qquad \text{(\# of corners)} \tag{6}$$

where $|\mathcal{C}|$ is the number of corners in the profile.

When operating in indoor environments, it is reasonable to assume that $z^\diamond$ has only a few corners. Therefore, we want to exploit this regularity assumption and the partial measurements y in (4) to reconstruct $z^\diamond$. Let us start from the noiseless case in which $\eta = \mathbf{0}$ in (4). In this case, a reasonable way to reconstruct the profile $z^\diamond$ is to solve the following optimization problem:

$$\min_z \|Dz\|_0 \quad \text{subject to } Az = y \qquad \text{(L0)}$$

which seeks the profile z that is consistent with the measurements (4) and contains the smallest number of corners. Unfortunately, problem (L0) is NP-hard due to the nonconvexity of the ℓ_0 (pseudo) norm. In this work we study the following relaxation of problem (L0):

$$\min_z \|Dz\|_1 \quad \text{subject to } Az = y \qquad (\mathrm{L1}_D)$$

which is a convex program (it can be indeed rephrased as a linear program), and can be solved efficiently in practice. Sections V provides conditions under which ($\mathrm{L1}_D$) recovers the solution of (L0). Problem ($\mathrm{L1}_D$) falls in the class of the cosparsity models in CS [33].

In the presence of bounded measurement noise (4), i.e., $\|\eta\|_\infty \leq \varepsilon$, the ℓ_1-minimization problem becomes:

$$\min_z \|Dz\|_1 \quad \text{subject to } \|Az - y\|_\infty \leq \varepsilon \qquad (\mathrm{L1}_D^\varepsilon)$$

Note that we assume that the infinity norm of the noise η is bounded, since this naturally reflects the sensor model in our robotic applications (i.e., bounded error in each laser beam). On the other hand, most CS literature considers the

[1] Note that x corresponds to the horizontal axis in Fig. 2(a), while the depth z is shown on the vertical axis in the figure.

ℓ_2 norm of the error to be bounded and thus obtains an optimization problem with the ℓ_2 in the constraint. The use of the ℓ_∞ norm as a constraint in ($\mathrm{L1}_D^\varepsilon$) resembles the *Dantzig selector* of Candes and Tao [12], with the main difference being the presence of the matrix D in the objective. We provide more in-depth discussions regarding the equivalence of the constrainted optimization problem ($\mathrm{L1}_D^\varepsilon$) and its uncontrained form. Interested reader may refer to Appendix B.

B. 3D Depth Reconstruction

In this section we discuss how to recover a 3D depth profile $Z^\diamond \in \mathbb{R}^{r\times c}$ (a depth map, as the one in Fig. 1(a)), using incomplete measurements. As in the 2D setup, we do not have direct access to $Z^\diamond$, but instead only (point-wise) $m \ll r \times c$ measurements in the form:

$$y_{i,j} = Z^\diamond_{i,j} + \eta_{i,j} \tag{7}$$

where $\eta_{i,j} \in \mathbb{R}$ represents measurement noise. Each measurement is a noisy sample of the depth of $Z^\diamond$ at pixel (i,j).

We assume that $Z^\diamond$ is sufficiently regular, which intuitively means that the depth profile contains many planar regions and few "edges". We define the edges as follows.

Definition 5 (Edge set). *Given a 3D profile $Z \in \mathbb{R}^{r\times c}$, the* vertical edge set $\mathcal{E}_V \subseteq \{2,\dots,r-1\}\times\{1,\dots,c\}$ *is the set of indices (i,j) such that $Z_{i-1,j} - 2Z_{i,j} + Z_{i+1,j} \neq 0$. The* horizontal edge set $\mathcal{E}_H \subseteq \{1,\dots,r\}\times\{2,\dots,c-1\}$ *is the set of indices (i,j) such that $Z_{i,j-1} - 2Z_{i,j} + Z_{i,j+1} \neq 0$. The* edge set $\mathcal{E}$ *is the union of the two sets: $\mathcal{E} \doteq \mathcal{E}_V \cup \mathcal{E}_H$.*

Intuitively, (i,j) is *not* in the edge set $\mathcal{E}$ if the 3×3 patch centered at (i,j) is planar, while $(i,j) \in \mathcal{E}$ otherwise. As in the 2D case we introduce 2nd-order difference operators D_V and D_H to compute the vertical differences $Z_{i,j-1} - 2Z_{i,j} + Z_{i,j+1}$ and the horizontal differences $Z_{i-1,j} - 2Z_{i,j} + Z_{i+1,j}$:

$$D_V Z^\diamond \in \mathbb{R}^{(r-2)\times c}, \qquad Z^\diamond D_H^\mathsf{T} \in \mathbb{R}^{r\times(c-2)} \tag{8}$$

where the matrices D_V and D_H are the same as the one defined (5), but with suitable dimensions; each entry of the matrix $D_V Z^\diamond$ contains the vertical (2nd-order) differences at a pixel, while $Z^\diamond D_H^\mathsf{T}$ collects the horizontal differences.

Following the same reasoning of the 2D case, we obtain the following ℓ_1-norm minimization

$$\begin{aligned} \min_Z \quad & \|\operatorname{vec}(D_V Z)\|_1 + \|\operatorname{vec}\big(Z D_H^\mathsf{T}\big)\|_1 \\ \text{subject to} \quad & Z_{i,j} = y_{i,j} \quad \text{for each measured pixel } (i,j) \end{aligned} \tag{9}$$

where $\operatorname{vec}(\cdot)$ denotes the (column-wise) vectorization of a matrix, and here we assume noiseless measurements. In the presence of measurement noise, the equality constraint in (9) is again replaced by an inequality $|Z_{i,j} - y_{i,j}| \leq \varepsilon$, $\forall (i,j)$, where ε is an upper bound on the pixel-wise noise $\eta_{i,j}$.

C. Reconciling 2D and 3D Depth Reconstruction

In this section we show that the 3D depth reconstruction problem (9) can be reformulated to be closer to its 2D counterpart ($\mathrm{L1}_D$), if we vectorize the depth profile (matrix Z). For a given profile $Z \in \mathbb{R}^{r\times c}$, we define the number of pixels $n \doteq r \times c$, and we call z the vectorized version of Z, i.e., $z \doteq \operatorname{vec}(Z) \in \mathbb{R}^n$. Using standard properties of the vectorization operator, we get

$$\begin{aligned} \operatorname{vec}(D_V Z) &= (\mathbf{I}_c \otimes D_V) z \\ \operatorname{vec}\big(Z D_H^\mathsf{T}\big) &= (D_H \otimes \mathbf{I}_r) z \\ Z_{i_j} = \operatorname{vec}\big(e_i^\mathsf{T} Z e_j\big) &= (e_i^\mathsf{T} \otimes e_j^\mathsf{T}) z \end{aligned} \tag{10}$$

where $\otimes$ is the Kronecker product, $\mathbf{I}_r$ is an identity matrix of size $r\times r$, and e_i is a vector which is zero everywhere except the i-th entry which is 1. Stacking all measurements (7) in a vector $y \in \mathbb{R}^m$ and using (10), problem (9) can be written succinctly as in ($\mathrm{L1}_\Delta$):

$$\min_z \|\Delta z\|_1 \quad \text{subject to } Az = y \tag{$\mathrm{L1}_\Delta$}$$

where the matrix $A \in \mathbb{R}^{m\times n}$ (stacking rows in the form $e_i^\mathsf{T} \otimes e_j^\mathsf{T}$) has the same structure of the sampling matrix introduced in Definition 3, and the "regularization" matrix Δ is:

$$\Delta \doteq \begin{bmatrix} \mathbf{I}_c \otimes D_V \\ D_H \otimes \mathbf{I}_r \end{bmatrix} \in \mathbb{R}^{2(n-r-c)\times n} \tag{11}$$

Note that ($\mathrm{L1}_\Delta$) is the same as ($\mathrm{L1}_D$), except for the fact that the matrix D in the objective is replaced with a larger matrix Δ. It is worth noticing that the matrix Δ is extremely sparse, with only 3 non-zero entries (1, -2, and 1) on each row in suitable (but not necessarily consecutive) positions.

In the presence of noise, we define an error vector $\eta \in \mathbb{R}^m$ which stacks the noise terms in (7) for each pixel (i,j), and assume pixel-wise bounded noise $\|\eta\|_\infty \leq \varepsilon$. The noisy 3D depth reconstruction problem then becomes:

$$\min_z \|\Delta z\|_1 \quad \text{subject to } \|Az - y\|_\infty \leq \varepsilon \tag{$\mathrm{L1}_\Delta^\varepsilon$}$$

Again, comparing ($\mathrm{L1}_D^\varepsilon$) and ($\mathrm{L1}_\Delta^\varepsilon$), it is clear that in 2D and 3D we solve the same optimization problem, with the only difference lying in the matrices D and Δ.

V. Analysis: Conditions for Exact Recovery and Error Bounds for Noiseless and Noisy Reconstructions

This section provides a comprehensive analysis on the quality of the depth profiles reconstructed by solving problems ($\mathrm{L1}_D$) and ($\mathrm{L1}_D^\varepsilon$) in the 2D case, and problems ($\mathrm{L1}_\Delta$) and ($\mathrm{L1}_\Delta^\varepsilon$) in 3D. In particular, Section V-A discusses *exact recovery* and provides the conditions on the depth measurements such that the full depth profile can be recovered exactly. Since these conditions are quite restrictive in practice (although we will discuss an interesting application to data compression in Section VII), Section V-B analyzes the reconstructed profiles under more general conditions. More specifically, we derive error bounds that quantify the distance between the ground truth depth profile and our reconstruction. Section V-C extends these error bounds to the case in which the depth measurements are noisy.

A. Sufficient Conditions for Exact Recovery

In this section we provide sufficient conditions under which the full depth profile can be reconstructed exactly from the given depth samples.

Recent results on cosparsity in compressive sensing provide sufficient conditions for exact recovery of a cosparse profile $z^\diamond$, from measurements $y = Az^\diamond$ (Proposition 6 below holds for a generic A). After presenting the result we discuss why this is not a condition roboticists would use.

Proposition 6 (Exact Recovery [45]). *Consider a vector $z^\diamond \in \mathbb{R}^n$ with D-support $\mathcal{I}$ and D-cosupport $\mathcal{J}$. Define $\bar{m} \doteq n - m$. Let $N \in \mathbb{R}^{\bar{m}\times n}$ be a matrix whose rows span the null space of the matrix A. Let $(\cdot)^\dagger$ denote the Moore-Penrose pseudoinverse of a matrix. If the following holds:*

$$C_{er} \doteq \|(N(D_{\mathcal{J}})^{\mathsf{T}})^\dagger N(D_{\mathcal{I}})^{\mathsf{T}})\|_{\infty\to\infty} < 1 \qquad (12)$$

then problem ($\mathrm{L1}_D$) *recovers $z^\diamond$ exactly.*

Despite the generality, Proposition 6 provides only an *algebraic* condition. In the depth estimation problem, *geometric* conditions that suggest the best sampling locations would be much more meaningful. Our contribution in this section is an interpretation of Proposition 6 in geometric terms.

We provide a first result for the 2D case. The proof is in Appendix C.

Proposition 7 (Exact Recovery of 2D depth profiles). *Let $z^\diamond \in \mathbb{R}^n$ be a 2D depth signal with corner set $\mathcal{C}$. Assuming noiseless measurements* (4)*, the following hold:*

(i) if the sampling set $\mathcal{M}$ is the union of the corner set and the first and last entries of $z^\diamond$, then $C_{er} = 1$;

(ii) if the sampling set $\mathcal{M}$ includes the corners and their neighbors (adjacent entries), then $C_{er} = 0$ and problem ($\mathrm{L1}_D$) *recovers $z^\diamond$ exactly.*

Proposition 7 implies that we can recover the original profile exactly, if we measure the neighborhood of each corner. An example that satisfies such condition is illustrated in Fig. 3(a). When we sample only the corners, however, Proposition 7 states that $C_{\text{er}} = 1$; in principle in this case one might still hope to recover the profile $z^\diamond$, since the condition $C_{\text{er}} < 1$ in Proposition 6 is only *sufficient* for exact recovery. But it turns out that in our problem one can find counterexamples with $C_{\text{er}} = 1$ in which ℓ_1-minimization fails to recover $z^\diamond$. A pictorial example is shown in Fig. 3(b), where we show an optimal solution $z^\star$ which differs from the true profile $z^\diamond$.

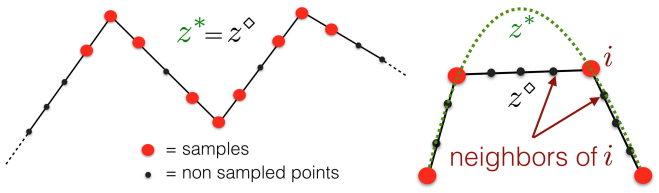

(a) (b)

Fig. 3. (a) Sampling corners and neighbors in 2D depth estimation guarantees exact recovery of the true profile. (b) Sampling only the corners does not guarantee, in general, that the solution of ℓ_1-minimization, namely $z^\star$, coincides with the true profile $z^\diamond$.

We derive a similar condition in 3D problems. The proof is given in Appendix D.

Proposition 8 (Exact Recovery of 3D depth profiles). *Let $Z^\diamond$ be a 3D depth signal with edge set $\mathcal{E}$. Assume noiseless measurements. If the sampling set $\mathcal{M}$ includes the edges and theirs (vertical and horizontal) neighbors (adjacent pixels), then $C_{er} = 0$, and* ($\mathrm{L1}_\Delta$) *recovers $vec(Z^\diamond)$ exactly.*

In the experimental section, we show that these initial results already unleash interesting applications. For instance, in stereo vision problems, we could locate the position of edges from the RGB images and recover the depth in a neighborhood of the edge pixels. Then, the complete depth profile can be recovered (at arbitrary resolution) via ($\mathrm{L1}_\Delta$).

B. Depth Reconstruction from Noiseless Samples

The exact recovery conditions of Proposition 7 and 8 are quite restrictive if we do not have prior knowledge of the position of the corners or edges. In this section we provide more powerful results that do not require sampling corners or edges. Empirically, we observe that when we do not sample edges, the optimization problems ($\mathrm{L1}_D$) and ($\mathrm{L1}_\Delta$) admit multiple solutions, i.e., multiple signals z attain the same optimal cost. The basic questions addressed in this section are: *which profiles are in the solution set$\mathcal{S}^\star$ of problems* ($\mathrm{L1}_D$) *and* ($\mathrm{L1}_\Delta$)*? Is the ground truth profile $z^\diamond$ among these optimal solutions? How far can an optimal solution be from the ground truth profile $z^\diamond$?* In order to answer these questions, in this section we derive optimality conditions for problems ($\mathrm{L1}_D$) and ($\mathrm{L1}_\Delta$), under the assumption that all mesurements are noise-free.

1) **Algebraic Optimality Conditions***:* In this section, we derive a general algebraic condition for a 2D signal (resp. 3D) to be in the solution set of ($\mathrm{L1}_D$) (resp. ($\mathrm{L1}_\Delta$)). Section V-B.2 and Section V-B.3 translate this algebraic condition into a geometric constraint on the curvature of the signals in the solution set.

Proposition 9 (2D Optimality). *Let A be the sampling matrix and $\mathcal{M}$ be the sample set. Given a profile $z \in \mathbb{R}^n$ which is feasible for* ($\mathrm{L1}_D$)*, z is a minimizer of* ($\mathrm{L1}_D$) *if and only if there exists a vector*
u such that

$$(D^{\mathsf{T}})_{\overline{\mathcal{M}}}\, u = \mathbf{0} \;\; \textit{and} \;\; u_{\mathcal{I}} = sign(Dz)_{\mathcal{I}} \;\; \textit{and} \;\; \|u\|_\infty \leq 1 \qquad (13)$$

where $\mathcal{I}$ is the D-support of z (i.e., the set of indices of the nonzero entries of Dz) and $\overline{\mathcal{M}}$ is the set of entries of z that we do not sample (i.e., the complement of $\mathcal{M}$).

The proof for both Proposition 9 is based on the subdifferential of the ℓ_1-minimization problem and is provided in Appendix E. An analogous result holds in 3D.

Corollary 10 (3D optimality). *A given profile Z is in the set of minimizers of* ($\mathrm{L1}_\Delta$) *if and only if the conditions of Proposition 9 hold, replacing Δ with D in eqs.* (13)*.*

We the skip the proof for Corollary 20 since the same mechanism for the proof of Proposition 9 can be applied.

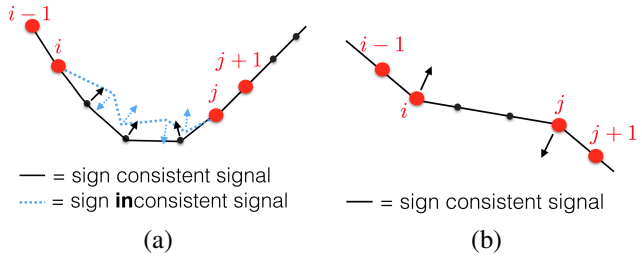

Fig. 4. Examples of sign-consistent (shown in black) and sign-inconsistent (shown in blue) 2D depth profile.

2) **Geometric Conditions for Optimality in 2D***:* In this section we derive *necessary and sufficient geometric conditions* for $z^\diamond$ to be in the solution set of ($\mathrm{L1}_D$). Using these findings we obtain two practical results: (i) an upper bound on how far any solution $z^\star$ of ($\mathrm{L1}_D$) can be from the ground truth signal $z^\diamond$; (ii) a general algorithm that recovers $z^\diamond$ even when the conditions of Proposition 7 fail.

To introduce our results, we need the following definition.

Definition 11 (2D Sign Consistency). *Let $s_k = sign(z_{k-1} - 2z_k + z_{k+1})$ (sign of the curvature at k). A 2D depth signal z is* sign consistent *if, for any two consecutive samples $i < j \in \mathcal{M}$, one of the two conditions holds:*

(i) no sign change: for any two integers k, h, with $i \leq k, h \leq j$, if $s_k \neq 0$ and $s_h \neq 0$, then $s_k = s_h$;

(ii) sign change only *at the boundary: for any integer k, with $i < k < j$, $s_k = 0$;*

This technical definition has a clear geometric interpretation. More specifically, a signal z is sign consistent, if its curvature does not change sign (i.e., it is either convex or concave) within each interval between consecutive samples. See Fig. 4 for examples of sign consistency, alongside with a counter-example.

In the following we show that any optimal solutions for problem ($\mathrm{L1}_D$) must be sign consistent. In order to simplify the analysis for Theorem 13 below, we assume that we pick pairs of consecutive samples (rather than individual, isolated samples). We formalize such notion of "twin samples" as follows.

Definition 12 (Twin samples). *A* twin sample *is a pair of consecutive samples, e.g. $(i, i+1)$ with $i \in \{1, \dots, n-1\}$.*

Theorem 13 (2D Sign Consistency $\Leftrightarrow$ Optimality). *Let z be a 2D signal which is feasible for problem ($\mathrm{L1}_D$). Assume that the sample set includes only twin samples and we sample the "boundary" of the profile, i.e., z_1, and z_n. Then, z is optimal for ($\mathrm{L1}_D$) if and only if it is sign consistent.*

The proof for Theorem 13 is in Appendix F. This theorem provides a tight geometric condition for a profile to be optimal. More specifically, a profile is optimal for problem ($\mathrm{L1}_D$) if it passes through the specified set of samples (i.e., it satisfies the constraint in ($\mathrm{L1}_D$)) and does not change curvature between consecutive samples. This result also provides insights into the conditions under which the ground truth signal will be among the minimizers of ($\mathrm{L1}_D$), and how one can bound the depth estimation error, as stated in the following proposition.

Proposition 14 (2D Recovery Error - noiseless samples). *Let $z^\diamond$ be the ground truth generating noiseless measurements* (4)*. Assume that we sample the boundary of $z^\diamond$ and the sample set includes a twin sample in each linear segment in $z^\diamond$. Then, $z^\diamond$ is in the set of minimizers of* ($\mathrm{L1}_D$)*. Moreover, denote with $\tilde{z}$ the naive solution obtained by connecting consecutive samples with a straight line (linear interpolation). Then, any optimal solution $z^\star$ lies between $z^\diamond$ and $\tilde{z}$, i.e., for any index $i \in \{1, \dots, n\}$, it holds that if $z_i^\diamond \leq \tilde{z}_i$ then $z_i^\diamond \leq z_i^\star \leq \tilde{z}_i$, and if $\tilde{z}_i \leq z_i^\diamond$ then $\tilde{z}_i \leq z_i^\star \leq z_i^\diamond$. Moreover, it holds*

$$\|z^\diamond - z^\star\|_\infty \leq \max_{i \in \mathcal{M}} d_i \cos(\theta_i) \tag{14}$$

where d_i is the distance between the sample i and the nearest corner in $z^\diamond$, while θ_i is the angle that the line connecting i with the nearest corner forms with the vertical.

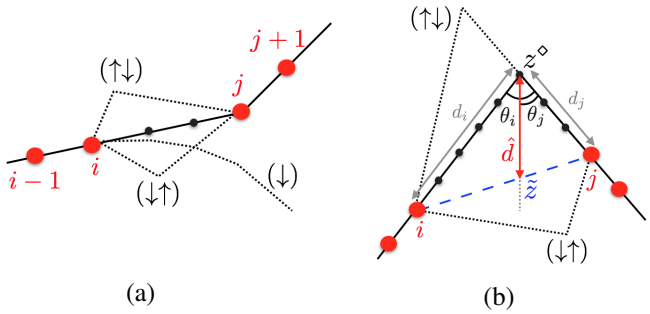

Fig. 5. (a) Region between a pair of twin samples, with a twin sample including a corner. (b) Region between a pair of twin samples, with twin samples not including corners.

A visualization of the parameters d_i and θ_i is given in Fig. 5(b). The proof is in Appendix G.

Proposition 14 provides two important results. First, it states that any optimal solution $z^\star$ (e.g., the dotted green line in Fig. 2(b)) should lie between the ground truth depth $z^\diamond$ (solid black line) and the naive solution $\tilde{z}$ (dashed blue line). In other words, any arbitrary set of twin samples has a corresponding *envelop* that contains all possible solutions. An example of such *envelop* is illustrated in Fig. 6(a). The width of this envelop bounds the maximum distance between any optimal solution and the ground truth, and hence such envelop provides a *point-wise quantification of the reconstruction error*. Second, Proposition 14 provides an upper bound on the overall reconstruction error in eq. (14). The inequality implies that the reconstruction error grows with

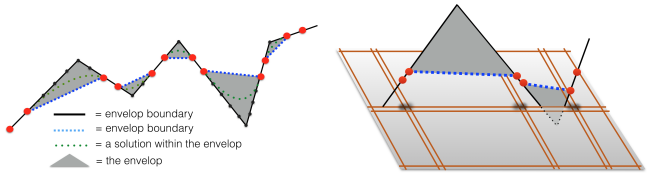

Fig. 6. (a) Given a set of noiseless samples, all possible optimal solutions in 2D are contained in the the envelop (triangles in gray). (b) A cross section of the envelop in 3D.

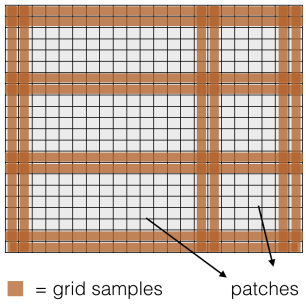

Fig. 7. Illustration of a grid sample set, along with 6 un-sampled patches in white.

parameter d_i, the distance between our samples and the corners. In addition, the error also increases if the parameter θ_i is small, meaning that the ground truth profiles are "pointy" and that there exist abrupt changes of slope between consecutive segments. An instance of such "pointy" signals is the second corner from right in Fig. 6(a).

We will further show in Section VI that Proposition 14 has algorithmic implications. Based on Proposition 14, we design an algorithm that exactly recovers a 2D profile, even when the sample set does not contain all corners. Before moving to these algorithmic aspects, let us consider the 3D case.

3) **Geometric Conditions for Optimality in 3D***:* In this section we provide a sufficient geometric condition for a 3D signal to be in the solution set of ($\mathrm{L1}_\Delta$). We start by introducing a specific sampling strategy (the analogous of the twin samples in 2D) to simplify the analysis.

Definition 15 (**Grid samples and Patches).** *Given a 3D signal $Z \in \mathbb{R}^{r\times c}$, a* grid sample set *includes pairs of consecutive rows and columns of Z, along with the boundaries (first and last two rows, first and last two columns). This sampling strategy divides the image in rectangular* patches*, i.e., sets of un-sampled pixels enclosed by row-samples and column-samples.*

Fig. 7 shows an example of grid samples and patches. If we have K patches and we denote the set of un-sampled pixels in patch i with $\overline{\mathcal{M}}_i$, then the union $\mathcal{M} \cup \{\overline{\mathcal{M}}_i\}_{i=1}^K$ includes all the pixels in the image. We can now extend the notion of sign consistency to the 3D case.

Definition 16 (**3D Sign Consistency).** *Let $Z \in \mathbb{R}^{r\times c}$ be a 3D depth signal. Let $\mathcal{M}$ be a grid sampling set and $\{\overline{\mathcal{M}}_i\}_{i=1}^K$ be the un-sampled patches. Let $Z_{[\overline{\mathcal{M}}_i]}$ be restriction of Z to its entries in $\overline{\mathcal{M}}_i$. Then, Z is called* 3D sign consistent *if for all $i = \{1,\dots,K\}$, the nonzero entries of sign$(vec(DZ_{[\overline{\mathcal{M}}_i]}))$ are all $+1$ or -1, and the nonzero entries of sign$(vec(Z_{[\overline{\mathcal{M}}_i]}D^\mathsf{T}))$ are all $+1$ or -1, where D is 2nd-order difference operator* (5) *of suitable dimension.*

Intuitively, 3D sign consistency indicates that the sign of the profile's curvature does not change, either horizontally or vertically, within each un-sampled patch. We now present a sufficient condition for $Z^\diamond$ to be in the solution set of ($\mathrm{L1}_\Delta$).

Theorem 17 (**3D Sign Consistency $\Rightarrow$ Optimality).** *Let $Z \in \mathbb{R}^{r\times c}$ be a 3D signal, feasible for problem* ($\mathrm{L1}_\Delta$)*. Assume the sample set $\mathcal{M}$ is a grid sample set. Then Z is in the set of minimizers of* ($\mathrm{L1}_\Delta$) *if it is 3D sign consistent.*

The proof is in Appendix H. Roughly speaking, if our grid sampling is "fine" enough to capture all changes in the sign of the curvature, then $Z^\diamond$ is among the solutions of ($\mathrm{L1}_\Delta$). Theorem 17 is weaker than Theorem 13, the 2D counterpart, since it states that sign consistency is only *sufficient*, but not necessary, for optimality.

Proposition 18 (**3D Recovery Error - noiseless samples).** *Let $Z^\diamond$ be the ground truth generating noiseless measurements* (4)*. Let $\mathcal{M}$ be a grid sampling set and assume $Z^\diamond$ to be* 3D sign consistent *with respect to $\mathcal{M}$. Let $\underline{Z}_{i,j}$ and $\bar{Z}_{i,j}$ be the lower and upper bound of the row-wise envelop. Then, given any optimal solution $Z^\star$ of* ($\mathrm{L1}_\Delta$)*, it holds that*

$$|Z^\diamond_{i,j} - Z^\star_{i,j}| \leq \max(|\underline{Z}_{i,j} - {Z^\star}_{i,j}|, |\bar{Z}_{i,j} - {Z^\star}_{i,j}|) \quad (15)$$

Despite the similarity to Proposition 14, the result in Proposition 18 is weaker. More specifically, Proposition 18 is based on the fact that we can compute an envelop *only* for the ground truth profile (but *not* for the optimal solutions, as in Proposition 14). Moreover, the estimation error bound in eq. (15) can be only computed *a posteriori*, i.e., after given an optimal solution $Z^\star$. Therefore, Proposition 18 does not provide a geometric understanding as clear as the bound in eq. (14).

C. Depth Reconstruction from Noisy Samples

In this section we analyze the depth reconstruction quality for the case where the measurements (4) are noisy. In other words, we now focus on problems ($\mathrm{L1}_D^\varepsilon$) and ($\mathrm{L1}_\Delta^\varepsilon$).

1) **Algebraic Optimality Conditions (noisy case)***:* In this section, we derive a general algebraic condition for a 2D signal (resp. 3D) to be in the solution set of ($\mathrm{L1}_D^\varepsilon$) (resp. ($\mathrm{L1}_\Delta^\varepsilon$)). These conditions are the equivalences of the optimality conditions for Section V-B.1 in the noisy case. In Section V-C.2 and Section V-C.3, we apply this algebraic condition to bound the depth reconstruction error.

Proposition 19 (**2D robust optimality).** *Let A be the sampling matrix, $\mathcal{M}$ be the sample set and y be the noisy measurements as in* (4)*, with $\|\eta\|_\infty \leq \varepsilon$ and $\varepsilon > 0$. Given a profile z which is feasible for* ($\mathrm{L1}_D^\varepsilon$)*, define the active set $\mathcal{A} \subset \mathcal{M}$ as follows*

$$\mathcal{A} = \{i \in \mathcal{M} : |y_i - z_i| = \varepsilon\}. \quad (16)$$

We also define its two subsets

$$\mathcal{A}_\uparrow = \{i \in \mathcal{M} : y_i - z_i = \varepsilon\},$$
$$\mathcal{A}_\downarrow = \{i \in \mathcal{M} : y_i - z_i = -\varepsilon\}. \quad (17)$$

Also denote $\bar{\mathcal{A}} = \mathcal{M} \setminus \mathcal{A}$. Then z is a minimizer of ($\mathrm{L1}_D^\varepsilon$) *if and only if there exists a vector u such that*

$$(D^\mathsf{T})_{\overline{\mathcal{M}}\cup\bar{\mathcal{A}}}\, u = \mathbf{0} \;\; \textit{and} \;\; u_\mathcal{I} = sign(Dz)_\mathcal{I} \;\; \textit{and} \;\; \|u\|_\infty \leq 1 \quad (18)$$
$$(D^\mathsf{T})_{\mathcal{A}_\uparrow}\, u \geq \mathbf{0} \;\; \textit{and} \;\; (D^\mathsf{T})_{\mathcal{A}_\downarrow}\, u \leq \mathbf{0} \quad (19)$$

where $\mathcal{I}$ is the D-support of z, and $\overline{\mathcal{M}}$ is the set of un-sampled entries in z (i.e., the complement of $\mathcal{M}$).

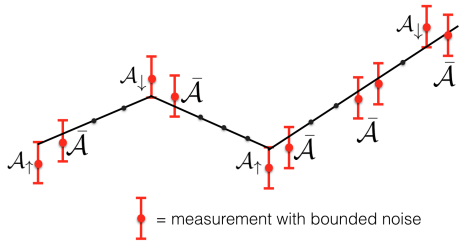

Fig. 8. A illustration of *active set* $\mathcal{A}$ of pixels in Proposition 19. The three corners, as well as the second sample from the right, are all within the active set. A measurement y_i is *active* if the reconstructed signal z hits the boundary of y_i's associated valid interval.

The proof is in Appendix J. A visual illustration of active set is in Fig. 8. We will provide some geometric insights on the algebraic conditions in Proposition 19 in the next two sections. Before moving on, we re-ensure that the robust optimality conditions straightforwardly extends to the 3D case.

Corollary 20 (3D robust optimality). *A given profile Z is in the set of minimizers of* $(\mathrm{L1}^{\varepsilon}_{\Delta})$ *if and only if the conditions of Proposition 19 hold, replacing Δ with D in eqs.* (18)-(19).

We the skip the proof for Corollary 20 since the same mechanism for the proof of Proposition 19 can be applied.

2) **Geometric Conditions for 2D Noisy Reconstruction***:* In this section we consider the $2D$ case and provide a geometric interpretation of the algebraic conditions in Proposition 19. The geometric interpretation follows from a basic observation, which enables us to relate the noisy case with our noiseless analysis of Section V-B.2. The observation is that if a profile satisfies the robust optimality conditions (18)-(19) then it also satisfies the noiseless optimality condition (13).

Theorem 21 (Robust optimality $\Rightarrow$ 2D Sign Consistency). *Let $z^\star$ be a 2D signal which is optimal for problem* $(\mathrm{L1}^{\varepsilon}_{D})$*, and assume that the sample set includes only twin samples and we sample the "boundary" of the profile, i.e., z_1, and z_n. Then, $z^\star$ is 2D sign consistent.*

We present a brief proof for Theorem 21 below.

Proof. Note that

$$(D^{\mathsf{T}})_{\overline{\mathcal{M}\cup\mathcal{A}}}\, u = \mathbf{0} \;\;\Rightarrow\;\; (D^{\mathsf{T}})_{\overline{\mathcal{M}}}\, u = \mathbf{0}.$$

In other words, condition (18) implies condition (13), which in turn is equivalent to sign consistency as per Theorem 13. Therefore, we come to the conclusion that any optimal solution of $(\mathrm{L1}^{\varepsilon}_{D})$ must be sign consistent, □

Theorem 21 will help establish error bounds on depth reconstruction. Now we introduce the formal definition of 2D sign consistent envelop.

Definition 22 (2D Sign Consistency Envelop). *Consider two consecutive pairs of twin samples $i, i+1$ and $j, j+1$. The 2D sign-consistent envelop contains the set of all 2D sign-consistent z, such that for all samples $k \in (i+1, j)$,*

$$\textit{(i)}\;\; \frac{z_k-(y_i-\varepsilon)}{x_k-x_i} \le \frac{(y_{i+1}+\varepsilon)-(y_i-\varepsilon)}{x_{i+1}-x_i}$$
$$\textit{(ii)}\;\; \frac{z_k-(y_i+\varepsilon)}{x_k-x_i} \ge \frac{(y_{i+1}-\varepsilon)-(y_i-\varepsilon)}{x_{i+1}-x_i}$$

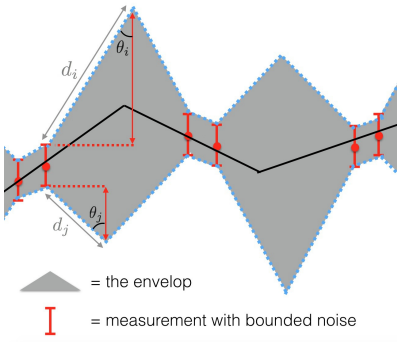

Fig. 9. An example of the envelop that contains all possible optimal solutions in 2D, given noisy samples.

$$\textit{(iii)}\;\; \frac{z_k-(y_{j+1}-\varepsilon)}{x_k-x_{j+1}} \ge \frac{(y_j+\varepsilon)-(y_{j+1}-\varepsilon)}{x_j-x_{j+1}}$$
$$\textit{(iv)}\;\; \frac{z_k-(y_{j+1}+\varepsilon)}{x_k-x_{j+1}} \le \frac{(y_j-\varepsilon)-(y_{j+1}+\varepsilon)}{x_j-x_{j+1}}$$

An example of such an envelop is illustrated in Fig. 9. Next we introduce a proposition that characterizes the error bounds of the optimal solutions.

Proposition 23 (2D Recovery Error - noisy samples). *Let $z^\diamond$ be the ground truth generating noisy measurements* (4)*. Assume that we sample the boundary of $z^\diamond$ and the sample set includes a twin sample in each linear segment in $z^\diamond$. Then, $z^\diamond$ is in the set of minimizers of* $(\mathrm{L1}^{\varepsilon}_{D})$*, and any optimal solution $z^\star$ also lies in the 2D sign consistent envelop. Moreover, any optimal solution $z^\star$ lies between $\underline{Z}_{i,j}$ and $\bar{Z}_{i,j}$, the lower and upper bound of the row-wise envelop. It holds that*

$$\|z^{\diamond}_{i,j} - {z^\star}_{i,j}\|_\infty \le \bar{Z}_{i,j} - \underline{Z}_{i,j} \le d_i\cos(\theta_i) + d_j\cos(\theta_j) + 2\varepsilon,$$

where $d_i, d_j, \cos(\theta_i), \cos(\theta_j)$ are all increasing functions of ε.

The proof is available at Appendix K.

3) **Geometric Conditions for 3D Noisy Reconstruction***:* In this section we characterize the error bounds of optimal solutions $Z^\star$ of $(\mathrm{L1}^{\varepsilon}_{\Delta})$ in the noisy case. The result is similar to its noiseless counterpart in Proposition 18.

Proposition 24 (3D Recovery Error - noisy samples). *Let $Z^\diamond$ be the ground truth generating noisy measurements* (4)*. Let $\mathcal{M}$ be a grid sampling set and assume $Z^\diamond$ to be* 3D sign consistent *with respect to $\mathcal{M}$. Let $\underline{Z}_{i,j}$ and $\bar{Z}_{i,j}$ be the lower and upper bound of the row-wise envelop. Then, given any optimal solution $Z^\star$ of* $(\mathrm{L1}^{\varepsilon}_{\Delta})$*, it holds that*

$$|Z^{\diamond}_{i,j} - Z^{\star}_{i,j}| \le \max(|\underline{Z}_{i,j} - {Z^\star}_{i,j}|, |\bar{Z}_{i,j} - {Z^\star}_{i,j}|) \qquad (20)$$

The proof for Proposition 24 applies the same mechanism as Proposition 18, so we skip the details.

VI. Algorithms and Fast Solvers

The formulations discussed so far, namely $(\mathrm{L1}_{D})$, $(\mathrm{L1}^{\varepsilon}_{D})$, $(\mathrm{L1}_{\Delta})$, $(\mathrm{L1}^{\varepsilon}_{\Delta})$, directly translate into algorithms: each optimization problem can be solved using standard convex programming routines and returns a depth profile.

We describe two algorithmic variants that further enhance the quality of the depth reconstruction in Section VI-A, and then we discuss a fast solver for the resulting ℓ_1 minimization problems in Section VI-B.

Algorithm 1: Exact recovery of 2D depth profiles.

input : Measurements y, and sample set $\mathcal{M}$, including boundary and twin samples
output: Original signal $z^\diamond$

/* solve ℓ_1-minimization */
1 create matrices A (Definition 3) and D (eq. (5)) ;
2 solve $(f^\star, z^\star) = \min_z \|Dz\|_1 \quad$ subject to $Az = y$;
/* populate a vector of signs $s \in \{-1, 0, +1\}^n$ */
3 **for** *consecutive twin samples* $(i-1, i)$, $(j, j+1)$ **do**
4 **foreach** $k \in \{i+1, \dots, j-1\}$ **do**
5 set $s_k = \text{sign}((z^\star_{j+1} - z^\star_j) - (z^\star_i - z^\star_{i-1}))$
/* recover $z^\diamond$ within the solution set */
6 $z^\diamond = \text{argmin}_z \; s^\mathsf{T} z \quad$ subject to $Az = y$ and $\|Dz\|_1 \leq f^\star$;
7 return $z^\diamond$.

A. Enhanced Recovery in 2D and 3D problems

In this section we describe other algorithmic variants for the 2D and 3D case. The first algorithm, given in Section VI-A.1, solves 2D problems and is inspired by Proposition 14. The second one, given in Section VI-A.2, solves 3D problems and is a variant of (L1_Δ).

1) **Enhanced Recovery in 2D problems***:* Proposition 14 dictates that any optimal solution of (L1_D) lies between the naive interpolation solution and the ground truth profile $z^\diamond$ (recall Fig. 2(b)). Algorithm 1 is based on a simple idea: on one hand, if the true profile is concave between two consecutive samples (*cf.* with the first corner in Fig. 2(b)), then we should increase the profile in this particular interval as much as possible (while still being within the optimal set of (L1_D) $z^\diamond$); on the other hand, if the shape is convex (second corner in Fig. 2(b)) we look for a solution with the minimum depth in the optimal set, since this is the closest to $z^\diamond$.

Algorithm 1 first solves problem (L1_D) and computes an optimal solution $z^\star$ and the corresponding optimal cost $f^\star$ (lines 1-2). Let us skip lines 3-5 for the moment and take a look at line 6: the constraints in this optimization problem include the same constraint of line 2 ($Az = y$), plus an additional constraint in line 2 ($\|Dz\|_1 \leq f^\star$) that restricts z to stay within the optimal solution set. Therefore, it only remains to design a new objective function that "encourages" a solution that is close to $z^\diamond$ while still being within this optimal set. To this end, we use a simple linear objective $s^\mathsf{T} z$, where $s \in \{0, \pm 1\}^n$ is a vector of coefficients, such that the objective function penalizes the entries of signal z at $s_k = +1$ and rewards entries corresponding to $s_k = -1$. More specifically, the procedure for choosing a proper coefficient s_k is as follows. For any consecutive pairs of twin samples $(i-1, i)$ and $(j, j+1)$, the algorithm looks at the slope difference between the second pair (i.e., $z^\star_{j+1} - z^\star_j$) and the first pair ($z^\star_i - z^\star_{i-1}$). If this difference is negative, then the function $z^\star$ is expected to be concave between the samples. In this case the sign s_k for any point k between the samples is set to -1. If the difference is positive, then the signs are set to $+1$. Otherwise the signs will be 0. We prove the following.

Corollary 25 (Exact Recovery of 2D profiles by Algorithm 1). *Under the assumptions of Proposition 14, Algorithm 1 recovers the 2D depth profile $z^\diamond$ exactly.*

The proof is in Appendix L. Although Algorithm 1 is designed for noiseless samples, in the experiments in Section VII-A we also test a noisy variant by substituting the constraints in lines 2 and 6 with $\|Az - y\|_\infty \leq \varepsilon$.

2) **Enhanced Recovery in 3D problems***:* In the formulations (L1_Δ) and ($\text{L1}^\varepsilon_\Delta$) we used the matrix Δ to encourage "flatness", or in other words, regularity of the depth profiles. In this section we discuss several variants of objective functions.

`L1` *objective function*: The first and simplest objective function is the sum of the second order derivatives along both the horizontal and vertical directions, and is denoted as `L1`. More specifically, the objective function is

$$\|\Delta z\|_1^{\texttt{L1}} = \|\operatorname{vec}(z * K_{xx})\|_1 + \|\operatorname{vec}(z * K_{yy})\|_1,$$

where z is the depth image to reconstruct, "$*$" is the convolution operation in image processing, and the convolution kernels K_{xx} and K_{yy} are defined as

$$K_{xx} = \begin{bmatrix} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad K_{yy} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & -2 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

`L1diag` *objective function*: The second objective function be denoted as `L1diag`. Compared with `L1`, the second objective function includes an additional second order derivative along the diagonal direction. This additional term is $\|\operatorname{vec}(z * K_{xy})\|_1$. Therefore, the objective function is

$$\|\Delta z\|_1^{\texttt{L1diag}} = \|\Delta z\|_1^{\texttt{L1}} + \|\operatorname{vec}(z * K_{xy})\|_1$$

The kernel K_{xy} for the diagonal term is as follows.

$$K_{xy} = \frac{1}{4} * \begin{bmatrix} -1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & -1 \end{bmatrix}$$

Note that with both `L1` and `L1diag`, the Δ matrix corresponds to a 2^{nd} order difference operator on the image plane (i.e., the pixel space), but with different kernels.

`L1cart` *objective function*: The third objective function, denoted as `L1cart`, applies the 2^{nd} order difference operator on the cartesian space (rather than on the pixel space). The objective function is defined as follows

$$\|\Delta z\|_1^{\texttt{L1cart}} = \|\operatorname{vec}\big(z * K_{xx}^{\text{L1-cart}}\big)\|_1 + \|\operatorname{vec}\big(z * K_{yy}^{\text{L1-cart}}\big)\|_1$$

Moreover, the convolution kernels for `L1cart` depend on the actual (x, y)-coordinates of the variables, and are defined as follows.

$$K_{xx}^{\text{cart}}(i,j) = \begin{bmatrix} 0 & \frac{1}{x_{i,j} - x_{i-1,j} + c} & 0 \\ 0 & -\left(\frac{1}{x_{i,j} - x_{i-1,j} + c} + \frac{1}{x_{i+1,j} - x_{i,j} + c}\right) & 0 \\ 0 & \frac{1}{x_{i+1,j} + c} & 0 \end{bmatrix}^T,$$

$$K_{yy}^{\text{cart}}(i,j) = \begin{bmatrix} 0 & \frac{1}{y_{i,j} - y_{i,j-1} + c} & 0 \\ 0 & \frac{1}{y_{i,j} - y_{i,j-1} + c} - \frac{1}{y_{i,j+1} - y_{i,j} + c} & 0 \\ 0 & \frac{1}{y_{i,j+1} - y_{i,j} + c} & 0 \end{bmatrix}$$

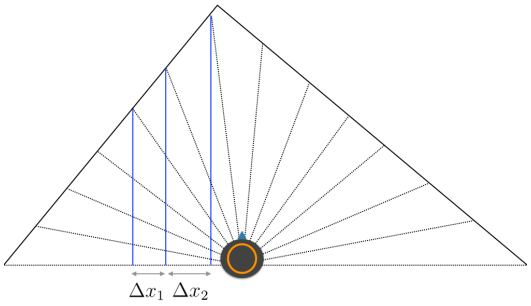

Fig. 10. A two-dimensional example illustrating the non-uniform spacing between measurements in cartisian coordinates, where $\Delta x_1 \neq \Delta x_2$ due to perspectie projection. This occurs when using LiDAR and cameras, and thus motivating the `L1cart` objective function.

This `L1cart` objective function is motivated by fact that the coordinates in the x, y dimensions have non-uniform spacing, as illustrated in Fig. 10. This phenomenon is due to the perspective projection of sensors such as a LiDAR and a camera. `L1cart` is expected to produce better reconstruction results, since it captures the actual geometry of the environment. In comparison, both `L1` and `L1diag` suffer from distortion due to perspective projection, and thus are can be viewed as simplified approximate of `L1cart`. However, additional technical issues arise with using `L1cart`: (1) `L1cart` requires a prior knowledge of all the (x, y) coordinates. These coordinates might come from, for example, a predetermined grid of query points where the depth of each point is of interest. (2) In order to avoid the ill-conditioned case with division by zero (e.g. when $x_{i,j} - x_{i-1,j} = 0$), we have to add a small positive constant c to all denominators.

In Section VII-C.1 we will have an empirical evaluation of these different objective functions.

B. Fast Solvers

All the formulations presented in this paper, as well as the algorithmic variants proposed in Section VI-A, rely on solving the optimization problems (L1_D), (L1_Δ), ($\text{L1}^\varepsilon_\Delta$), and ($\text{L1}^\varepsilon_D$) efficiently. Despite the convexity of these problems, off-the-shelf solvers based on interior point methods tend to be slow and do not scale to large problems. For example, recall that in the 3D case, the number of unknown variables in our problems is equal to the number of un-sampled pixels in the depth map, and hence can easily go beyond 10^4. Indeed, in the experiments in Section VII-C.2 we show that off-the-self solvers such as CVX/MOSEK [28] are quite slow and practically unusable for 3D profiles larger than 100×100 pixels.

For these reasons, we provide a more efficient first-order method to solve these minimization problems. This solver is a variant of NESTA, an algorithm for fast ℓ_1 minimization recently developed by Becker *et al.* [4] with basis on Nesterov's method for nonsmooth optimization [48], [49]. We tailor NESTA to our specific optimization problems with ℓ_∞ norm constraints, instead of the original ℓ_2 norm used in [4]. In this section we focus on the 2D problem (L1^ε_D), since the algorithm is identical in the $3D$ case (with the only exception that the matrix Δ is used in place of D).

Algorithm 2: NESTA for solving (L1^ε_D)

input : Measurements y, sampling matrix A, noise level ε, initial point z_0, desired final smoothing parameter value μ_f, maximum Nesterov iterations K, continuation steps T, stopping criterion τ

output: Approximate solution $z^{(K)}$ for (L1^ε_D)

initialize μ_0 ;

compute $\gamma = (\mu_f - \mu_0)^{1/T}$;

`/* Continuation steps */`

for $t = 1 : T$ **do**

  set $\mu = \mu_{t-1}$;

  `/* apply Nesterov's algorithm with updated` μ `and` $z^{(0)}$ `*/`

  for $k = 0 : K$ **do**

    compute $\nabla f_\mu(z^{(k)})$;

    set $\alpha_k = \frac{1}{2(k+1)}$ and $\tau_k = \frac{2}{k+3}$;

    `/* solve: */`

    $x^{(k)} = \operatorname{argmin}_z \frac{L_\mu}{2}\|z - z^{(k)}\|_2^2 + \langle \nabla f_\mu(z^{(k)}), z - z^{(k)} \rangle$

      subject to $\|Az - y\|_\infty \leq \varepsilon$;

    `/* solve: */`

    $y^{(k)} = \operatorname{argmin}_z \frac{L_\mu}{2}\|z - z^{(k)}\|_2^2 + \sum_{i=0}^{k} \alpha_i \langle \nabla f_\mu(z^{(i)}), z - z^{(k)} \rangle$

      subject to $\|Az - y\|_\infty \leq \varepsilon$;

    `/* update z */`

    $z^{(k)} = \tau_k y^{(k)} + (1 - \tau_k) x_k$;

    `/* break loop to speed up, if stopping criterion is met */`

    if $\|z^{(k)} - z^{(k-1)}\|_\infty < \tau$ **then**

      break loop

  `/* decrease the value of` μ `*/`

  set $\mu_t = \gamma \mu_{t-1}$;

  set $z^{(0)} = z^{(K)}$;

return $z^{(K)}$.

We provide an outline of this algorithm for the optimization problem (L1^ε_D) in this section, and leave the technical computations to Appendix M. In its most generality, NESTA aims to solve

$$\min_z \ f_\mu(z) \quad \text{subject to } z \in \mathcal{Q} \tag{21}$$

where $f_\mu(z)$ is a smooth approximation of a convex but nonsmooth function $f(z)$ and $\mathcal{Q}$ is a convex set. In our problem (L1^ε_D), we have $f(z) = \|Dz\|_1$ and $\mathcal{Q} = \{z : \|Az - y\|_\infty \leq \varepsilon\}$. An equivalent formulation for $f(z)$ is given in terms of the dual ℓ_∞-ball

$$f(z) = \max_{u : \|u\|_\infty \leq 1} \langle u, Dz \rangle.$$

Then a convenient choice for $f_\mu(z)$ is

$$f_\mu(z) = \max_{u : \|u\|_\infty \leq 1} \langle u, Dz \rangle - \mu \frac{\|u\|_2^2}{2}. \tag{22}$$

The function $f_\mu(z)$ is differentiable and its gradient is Lipschitz with constant L_μ, see Appendix M for a definition and the explicit value. Instead of minimizing $f(z)$ with nonsmooth optimization techniques which are generally slow, Nesterov proposes to use an improved gradient method to minimize $f_\mu(z)$ with an optimal convergence rate of $\mathcal{O}(1/K^2)$ in the number of iterations K. The corresponding

steps of Nesterov's method are given in lines 5-11 of the proposed Algorithm 2. Closed form formulas for $\nabla f_\mu(z)$ and the solutions $x^{(k)}$ and $y^{(k)}$ (lines 8-10) are provided in Appendix M. The choice of μ in (22) reflects a trade-off between the speed of convergence (the convergence rate is proportional to μ) and the accuracy of the smoothed approximation f_μ, which consequently determines the NESTA's overall accuracy. According to experiments in [4], decreasing μ by a factor of 10 gives about 1 additional digit of accuracy on the optimal value. NESTA also involves a continuation step of updating the smoothing parameter μ in order to substantially accelerate the algorithm, see lines 3-14. This step solves a succesion of problems (21) with decreasing smoothing parameters from an initial value μ_0 to the desired value μ_f. The number T of continuation steps is of choice and it determines the step size $\gamma < 1$ in line 2. Solving (21) is faster with larger μ and using intermediate solutions as a warm start for the next problem increases the overall speed significantly. We use the naive interpolation solution $\tilde{z}$ as an initial point $z^{(0)}$ for the algorithm. The initial value $\mu_0 = \|D^\mathsf{T} y\|_\infty$.

Note that the choice $\varepsilon = 0$ in Algorithm 2 implies the noiseless problem ($\mathrm{L1}_D$). This only affects the solutions $x^{(k)}$ and $y^{(k)}$. Similarly, Algorithm 2 can be used for the problems ($\mathrm{L1}_\Delta$) and ($\mathrm{L1}^\varepsilon_\Delta$) with the matrix D replaced with Δ in the definition of $f(z)$. As discussed earlier, the choice of a nonzero μ_f in NESTA will result in an approximate solution to the optimal solution of ($\mathrm{L1}^\varepsilon_D$). Consequently, NESTA produces slightly less accurate solutions to the problems at hand, but is much faster than CVX. A detailed comparison of the performances of two solvers is provided in our simulations Section VII-C.2.

VII. Experiments

In this section we verify our theoretical derivations with experiments on synthetic, simulated and real data. Empirical evidence shows that our recovery techniques perform very well in practice, in both 2D and 3D environments. Our algorithm is also much more robust to noise than naive linear interpolation. A number of applications are discussed, including 2D mapping, dense depth reconstruction from sparse measurements, data compression applied to bandwidth-limited robot-server communication, and super-resolution of depth images.

For all following tests, we evaluate the accuracy of reconstruction by the average pixel-wise depth error in meters, defined as $\frac{1}{n}\|z^\star - z^\diamond\|_1$, where $z^\diamond$ is the ground truth and $z^\star$ is the reconstruction.

A. 2D Sparse Reconstruction

In this section, as a sanity check of our algorithm in the simplest settings, we apply our algorithm to reconstruct 2D depth profiles (e.g., the data returned by a 2D laser scanner).

We create a synthetic dataset which contains random piecewise linear depth profiles of size $n = 2000$, with given a fixed number of corners. We also create a simulated dataset on the *Stage* simulator [26], where a robot equipped with a laser scanner with only 10 beams moves in a 2D scenario. This dataset includes both odometry information and 2D laser measurements.

When possible, we compare three different reconstruction algorithms, including the `naive` linear interpolation (Matlab's command `intep1`), the `L1` algorithm (see Section VI-A.2) and the two-stage Algorithm 1. More specifically, the results produced by ($\mathrm{L1}_D$) and ($\mathrm{L1}_\Delta$) are denoted with `L1`. The estimate produced by the noisy problems ($\mathrm{L1}^\varepsilon_D$) and ($\mathrm{L1}^\varepsilon_\Delta$) are denoted with $\mathtt{L1}(\varepsilon = \cdot)$, where we specify the noise level in the parenthesis. Since the number of variables is small, we use the CVX solver for a higher accuracy in all 2D experiments.

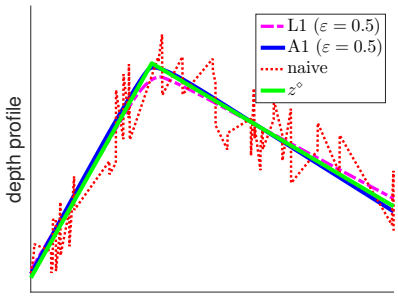

Fig. 11. An example of the synthetic 2D profiles (with only one corner).

1) Examples of 2D Reconstruction: An example of the synthetic 2D profiles (with only one corner) is shown in Fig. 11. The green plot is the ground truth signal, while the others are reconstructed depth profiles from sparse and noisy measurements using the three different algorithms.

From this figure it is clear that `naive` linearly interpolates the samples, hence even when measuring all depth data, it still produces a jagged line. It is easy to show that when measurement noise is uniformly distributed in $[-\varepsilon, +\varepsilon]$ (as in our tests), the error converges to $\varepsilon/2$ for increasing number of samples. On the other hand, `L1` and `A1` correctly smooths the noise out.

2) Statistics for 2D Reconstruction: Fig. 12(a)-(c) show how the reconstruction quality changes against the number of corners in the ground truth signal (i.e., as the signal gets less sparse), the percentage of measurements, and the noise level. Each data point in the plot is generated by averaging over 50 runs.

Fig. 12(a) compares `naive`, `L1`, and the estimate of Algorithm 1 (label: `A1`). These results consider noiseless measurements and sample set including a twin sample in each linear region (these are the assumptions of Proposition 14); reconstruction errors are shown for profiles with increasing number of corners. As predicted by Corollary 25, `A1` recovers the original signal exactly (zero error). `naive` has large errors, while the `L1` estimate falls between the two. Fig. 12(b) considers a more realistic situation: since in practice we do not know where the corners are, in this case we uniformly sample depth measurements and we consider noisy measurements with $\varepsilon = 0.1$m. As the percentage goes to 1 (100%), we sample all entries of the depth profile. We consider profiles with 3 corners in this test. The figure shows

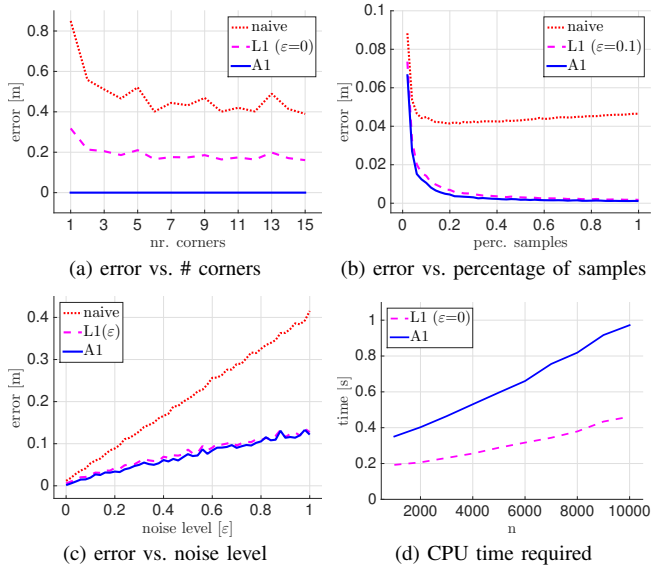

(a) error vs. # corners (b) error vs. percentage of samples

(c) error vs. noise level (d) CPU time required

Fig. 12. (a) Recovery errors for `A1`, `L1`, and `naive`. A twin sample is acquired in each linear segment and measurements are noiseless. (b) Recovery errors from uniformly sampled noisy measurements, with increasing number of samples. (c) Recovery errors from uniformly sampled noisy measurements (5% of the depth data points) and increasing noise level. (d) CPU time required to solve `L1` and `A1` for increasing size of the 2D depth signal (n).

that for increasing number of samples, our approach largely outperforms the `naive` approach. `A1` improves over `L1` even in presence of noise, while the improvement is not as substantial as in the noiseless case of Fig. 12(a). Fig. 12(b) also shows that the error committed by `naive` does not improve when adding more samples. This can be understood from the example of Fig. 11.

Fig. 12(c) considers a fixed amount of samples (5%) and tests three approaches for increasing measurement noise ε. Our techniques (`L1`, `A1`), are very resilient to noise and degrade gracefully even in presence of large noise (e.g., $\varepsilon = 1$m).

Fig. 12(d) shows the CPU times required by `L1` and `A1` on 2D reconstruction problems using `cvx` solver on MATLAB.

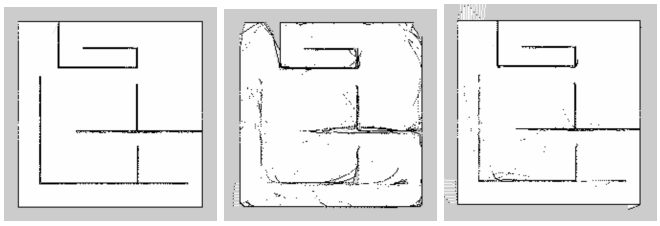

(a) full scan (b) `naive` (c) `A1`

Fig. 13. (a) Mapping using a conventional laser scanner, (b-c) Mapping using scans reconstructed from 10 measurements by `naive` and `A1`.

3) 2D mapping from sparse measurements: This section applies our approach to a 2D mapping problem with sparse measurements. We use the *Stage* simulator to simulate a robot equipped with a laser scanner with only 10 beams, moving in a 2D scenario. The robot is in charge of mapping the scenario; we assume the trajectory to be given. Our approach works as follows: we feed the 10 samples measured by our "sparse laser" to algorithm `A1`; `A1` returns a full scan (covering 180 degrees with 180 scans in our tests), which we feed to a standard mapping routine (we use `gmapping` [29] in our tests).

Fig. 13 compares the occupancy grid map produced by a standard mapping algorithm based on a conventional laser scan, against the occupancy grid map reconstructed from our 10-beam laser. Fig. 13(c) shows that we are able to do a fairly accurate reconstruction from very partial information. As a technical remark, we note that, before feeding the 10 laser measurements to `A1`, we express the corresponding points in Cartesian coordinates: the original data given by the laser is in polar coordinates and piecewise linear signals do not remain piecewise linear in polar coordinates (see Fig. 10).

B. 3D Datasets

In this section we introduce the datasets we use for sections to follow on 3D depth map reconstruction. In order to have a ground truth signal for proper evaluation of the reconstruction algorithms, we collected several datasets with commonly-used, regular-sized depth sensors (including a *Kinect* and a *ZED* stereo camera) and use the heavily downsampled depth image as our "sparse" depth measurements, instead of using actual sparse and low-resolution depth sensors.

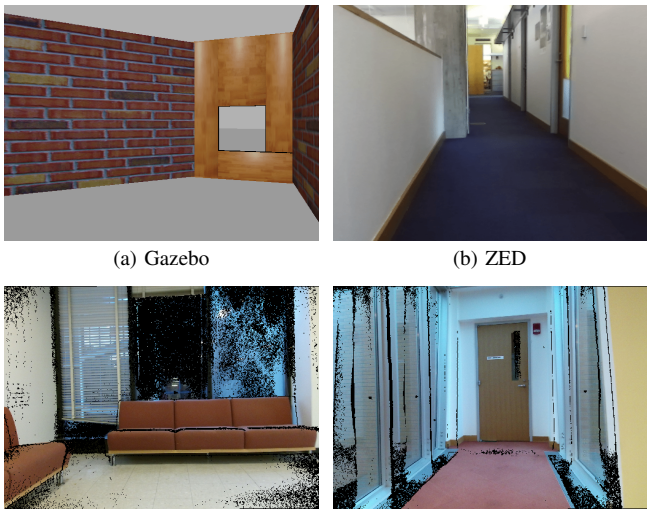

(a) Gazebo (b) ZED

(c) Kinect (d) Kinect

Fig. 14. (a) *Gazebo* Simulated data. (b) *ZED* Stereo data. (c)-(d) *Kinect* data.

The datasets that we collected include a dataset of synthetic, piecewise linear depth images (denoted as `PL`), a simulated dataset from the *Gazebo* simulator [35] (denoted as `Gazebo`), a stereo dataset from a *ZED* camera (denoted as `ZED`), and 8 datasets from a *Kinect* camera (denoted as `K1` to `K8`). More specifically, `Gazebo` contain 20 full depth and RGB images rendered in an office-like environment from the *Gazebo* simulator. `ZED` includes 1000 full disparity and RBG images, collected from a *ZED* stereo camera mounted on a dolly moving in the Laboratory of Information and Decision Systems (LIDS) at MIT. `K1` to `K8` contain odometry

information, as well as depth and RGB images, collected from a *Kinect* sensor mounted on a dolly with optical wheel odometers moving in 8 different locations at MIT, including tunnels, offices and corridors.

C. Evaluation of Objective Functions and Solvers

In Section IV we introduced both the exact reconstruction problem L1_Δ and the noisy reconstruction problem $\text{L1}_\Delta^\varepsilon$ in 3D. There exist several variants of the problem with slightly different regularization matrices Δ (and thus different objective functions to minimize), each with their own strength and weakness, as discussed in Section VI-A.2. There also exist several different solvers. In this section, we will have an empirical evaluation of these objective functions and solvers.

1) Objective Functions: In this section we compare three objective functions proposed in Section VI-A.2 for the exact reconstruction problem L1_Δ. We use the CVX/MOSEK [28] solver on MATLAB in this section for the highest accuracy.

We compare the reconstruction accuracy of the three different objective functions in Fig. 15 on 10 different datasets, including `PL`, `ZED`, and `K1`-`K8`. The error bars show the average reconstruction error for each on the objective functions (`L1`, `L1diag` and `L1cart`) on a dataset. The depth measurements are sampled from a grid (4% samples, i.e., one measurement from each 5-by-5 image patch).

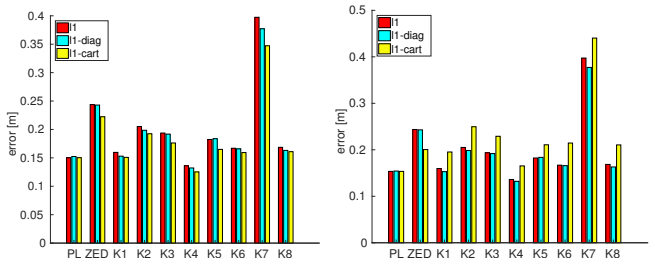

(a) (b)

Fig. 15. Comparison between three different cost functions, i.e., `L1` (without the diagonal terms in the cost function), `L1diag` (with the diagonal terms) and `L1cart` (use actual x, y coordinates to create the derivative matrix), on 10 different datasets. (a) constant c=0.1m (b) constant c=0.01m

Recall the discussion in Section VI-A.2 that a parameter c is inserted into the denominator when computing the second-order derivative, such that the ill-conditioned case of division by zero can be avoided. From Fig. 15 it is clear that `L1cart` depends heavily on the parameter c and degrades significantly with small value of c. Therefore, it is in general not reliable. The `L1diag` cost objective, on the other, performs better than `L1` consistently across all datasets and is parameter-free. Therefore, for the rest of the paper we will use `L1diag` for our all tests.

2) Solvers: In Section VI-B we discuss two different solvers for the ℓ_1-minimization problems. One of the solvers is CVX/MOSEK [28], a standard optimization toolbox for convex problems in MATLAB; and the other solver is NESTA [4], a first-order method for sparse recovery problems. In this section, we empirically compare the two solvers on their accuracy and speed.

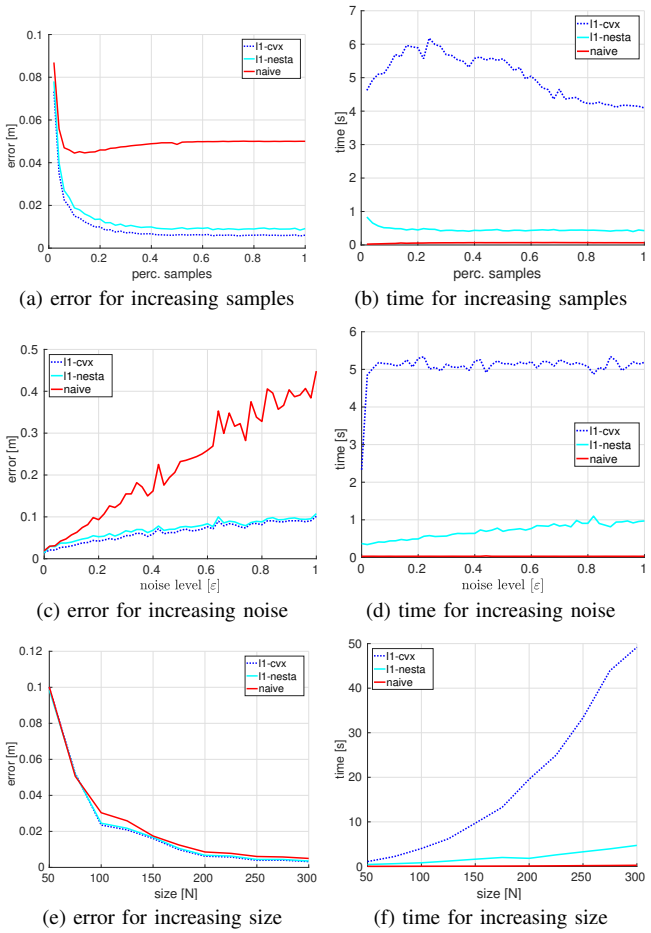

(a) error for increasing samples (b) time for increasing samples

(c) error for increasing noise (d) time for increasing noise

(e) error for increasing size (f) time for increasing size

Fig. 16. Comparison between NESTA ($\mu = 0.01$) and CVX/MOSEK. NESTA achieves almost exactly the same level of reconstruction error, while offering significant speedup.

We generate Fig. 16 by running experiments on the piecewise linear synthetic datasets `PL`. Each N-by-N depth image is generated randomly with a fixed number (3) of corners ($N = 100$ unless otherwise specified). All depth measurements are uniform-random samples from the ground truth signal, and the 4 immediate neighbors are also added into the sample sets. Each data point on the plot is averaged out over 50 trials, using both CVX/MOSEK and NESTA.

Fig. 16(a)-(b) show the reconstruction error and speed of the two solvers against increasing percentage of samples. Depth measurements are disturbed with noise with $\varepsilon = 10$ (specifically, the noise is uniformly random in the range of $[-\varepsilon, \varepsilon]$ and is added to the measurements in an entry-wise manner). Similarly, Fig. 16(c)-(d) show the reconstruction against increasing noise level. Fig. 16(e)-(f) show that reconstruction against increasing image size N (each random image is of size N-by-N) with 5% of samples, without adding noise. It is clear from Fig. 16 that the error curves for CVX/MOSEK and NESTA are almost indistinguishable, implying that they produce reconstructions of similar qualities. However, the NESTA solver has a speed up of 3x - 10x, depending on the exact setup of the reconstruction problem.

Given the significant advantage of NESTA over CVX/MOSEK, we will use NESTA for the rest of the paper. Note that NESTA's

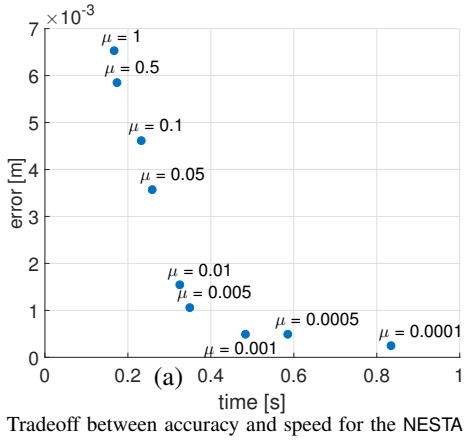

Fig. 17. Tradeoff between accuracy and speed for the NESTA solver with different parameter values μ. As μ decreases, the NESTA solver produces a more accurate solution at the cost of higher computational time.

accuracy is determined by the approximation factor μ, and in Fig. 17 we show the trade-off between reconstruction error and computational time for different parameters μ. Each data point is generated by averaging over 50 tests on piecewise linear synthetic data of size 100-by-100, with 5% uniformly random depth samples (and their immediate neighbors) each. For a balance of accuracy and speed, we will fix $\mu = 0.001$, the "elbow" point, for all following experiments.

D. Single-Frame 3D Sparse Reconstruction

In this section we demonstrate how our algorithms can reconstruct single-frame, complete 3D depth profiles from sparse and incomplete depth measurements (e.g., low-resolution laser scanner or sparse stereo correspondences). We evaluate the `L1diag` reconstruction by comparing its accuracy against the baseline approach, i.e., the `naive` linear interpolation. The `L1diag` problem is solved with NESTA with a fixed approximate parameter $\mu = 0.001$. We show examples of the reconstruction on different datasets, as well as complete error statistics with different percentages of samples and noise levels.

1) Examples: Here we show reconstruction examples based on sparse depth measurements from both `Gazebo` and `K1` datasets.

Fig. 18 (a)-(c) shows an example of `Gazebo` simulated dataset with uniformly random depth measurements (a sampling strategy that poses more serious challenges to the reconstruction algorithm) and the reconstructed full depth image based on these samples. The reconstructed depth image reflects the true geometry of the environment.

Fig. 18 (d)-(e) shows an example of the `K1` dataset, where all depth measurements fall on a regular grid. This sampling strategy resembles the output of a low-resolution depth sensor. Note that even though only a total number of 42 measurements measurements is available, the reconstructed depth image still correctly identifies the corridor and the walls.

2) Error Statistics: In this section we rigorously benchmark the performance of our algorithms in terms of both the reconstruction accuracy and the robustness to measurement noise.

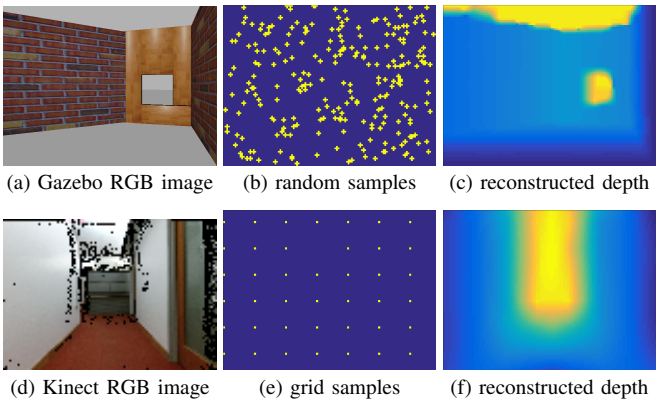

(a) Gazebo RGB image (b) random samples (c) reconstructed depth

(d) Kinect RGB image (e) grid samples (f) reconstructed depth

Fig. 18. The first row is an example of sparse depth reconstruction on `Gazebo` simulated data: (a) RGB image, (b) uniformly drawn sparse samples, and (c) reconstruction using `L1diag`. The second row is an example on Kinect `K1` data: (d) RGB image, (e) sparse samples on a grid, and (f) reconstruction using `L1diag`.

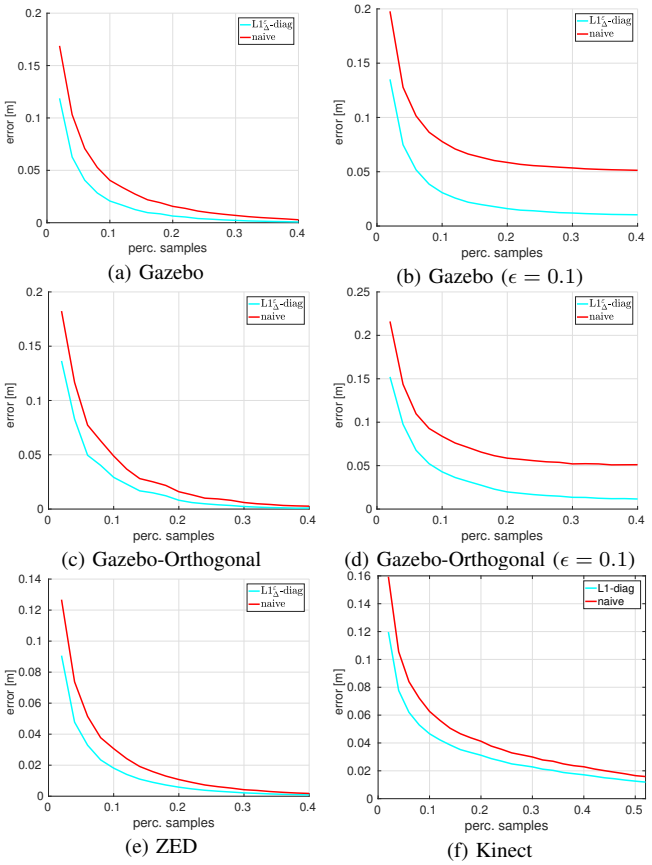

Fig. 19. Reconstruction errors for increasing percentage of uniform samples, and for different datasets. (a-c) reconstruction from noiseless samples, (d) reconstruction with measurements noise bounded by $\varepsilon = 0.1$m.

Fig. 19 depicts the reconstruction errors for increasing percentages of uniformly-random samples on different datasets. Fig. 19(a) and (c) show reconstruction from noiseless samples on the `Gazebo` simulated datasets, while Fig. 19(b) and (d) show the same plots except with additional pixel-wise independent Gaussian measurement noise $\varepsilon = 0.1$. Fig. 19(e)-(f) show the experimental results on the `ZED` stereo dataset and `K1` dataset. No additional noise is added to these two datasets, since the raw data from the sensors already contain real sensor noise.

From the figures it is clear that our approach consistently outperforms the baseline (i.e., `naive` linear interpolation) in both the noiseless and noisy settings and across different datasets. The gap between `L1diag` and `naive` widens as the number of samples increases under the noisy settings, which demonstrates that our approach is much more noise-resilient than the `naive` linear interpolation. Under the noiseless setup, the gap shrinks as the percentage of samples converges to 100%, since there's no loss of information. However, `L1diag` still produces significantly more accurate reconstruction (a 20% - 50% reduction of errors compared with `naive`) when operating at the below-20%-samples regime.

E. Multi-Frame 3D Sparse Reconstruction

In the previous Section VII-D we focus on sparse depth reconstruction with one single frame of data. However, when odometry information is available (e.g., from a wheel odometer equipped on ground vehicles, or from inertial measurement units on aerial robots), it is possible to combine sparse measurements across multiple frames over a time window in order to improve the reconstruction. More specifically, for every past frame in a given time window, we computate its associated sensor position (and pose) relative to the current frame by accumulating the odometry data. With this information, we project the set of measurements acquired at every past frames to the current frame. This way, we accumulate all measurements from the past time window. We reconstruct the dense depth image based on this set of measurements. Note that we assume the environment is static and the odometry is accurate within a short time window.

In this section, we demonstrate this idea with the `K1`-`K8` datasets, where the odometry information is available and can be considered reliable over a short temporal window.

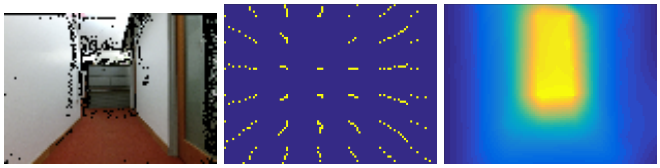

(a) environment (b) samples collected from 10 frames (c) reconstructed depth

Fig. 20. (a) registered RGB image from the Kinect sensor (b) depth samples collected across multiple time frames (c) reconstruction depth image. With more samples collected from several frames, the temporal reconstruction is more accurate (e.g., sharper at edges) than reconstruction using only one single frame.

1) Examples: Fig. 20 shows an example of multi-frame reconstruction. Fig. 20(a) is an RGB image (registered with the depth image, and thus some missing points). Fig. 20(b) shows the color-coded depth measurements collected over a temporal window across 10 frames. With odometry information, measurements collected at different time frames (and thus different locations and orientations) are projected to the last frame. Fig. 20(c) demonstrates the reconstructed depth image from all collected measurements. When compared against Fig. 18(e), it can be observed that with additional measurements the reconstructed depth image has sharper edges.

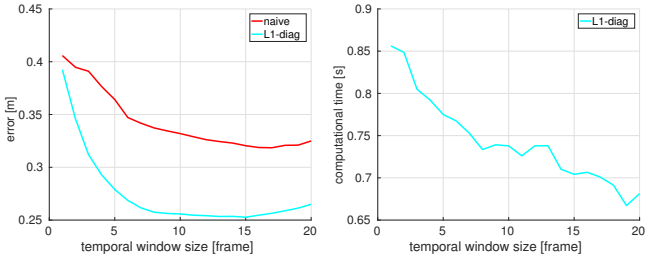

(a) (b)

Fig. 21. The two figures show error and timing statistics against increasing horizon size for temporal reconstruction. With more samples collected, the accuracy increases and the computational decreases (due to more constraints in the optimization problem). However, when the horizon size is very large, the error rises due to accumulated odometry error.

2) Statistics: Fig. 21 shows statistically the performance of multi-frame reconstruction over a temporal window, compared against the baseline `naive` linear interpolation. We run this simulation on `K1`, one of the Kinect datasets, and each data point on the plot is averaged over 910 images. From each frame only 18 depth measurements are collected, and the samples fall on a grid.

Fig. 21(a) contains error plots for `naive` and `L1diag`. It can be observed that `L1diag` outperforms `naive` consistently for different temporal window sizes. The best performance is achieved at window size of 15 with an average reconstruction error of 25cm, while the averaged error for the single-frame reconstruction (i.e., when the window size is 1) is 39cm, which is 56% higher.

Note that both error curves are convex, in the sense that the lowest error is achieved at window size of 15 and the error rises again with increasing length of temporal horizon. This phenomena can be attributed to the odometer drift (i.e., accumulation of odometry errors) over time. More specifically, the drift leads to incorrect point correspondence between frames, and thus point measurements in the past frames are projected to wrong pixel locations, resulting in overall performance degradation.

Fig. 21(b) shows the computational time for our algorithm `L1diag` using NESTA. The runtime decreases almost linearly with the length of temporal horizon. This is due to reduction of search space in the optimization problem thanks to additional constraints (measurements).

F. Data Compression

Another major application of our proposed algorithm lies in bandwidth-limited robot-server communication. Instead of having to transmit the entire depth image from the robot to a remote server, we show empirically that it is sufficient to send only depth data along the edges (extracted either from the RGB images or the depth images) and reconstruct the entire depth images on the server side. This way, we significantly reduce the required communication bandwidth. We measure the data compression ratio by *data rate saving*, defined as

$$\text{data rate saving} = 1 - \frac{\text{compressed data rate}}{\text{uncompressed data rate}}.$$

We demonstrate the compression technique on both the `Gazebo` and `ZED` datasets. We separate the discussions regarding edge extraction from RGB images and from the depth images. This is due to the fact that they have different pros and cons, and thus can be applied in different scenarios.

(a) RGB image (b) depth edges (c) reconstruction

Fig. 22. An example of depth image compression with data only along depth images: (a) RBG image, (b) edges in the depth image, (c) `L1diag` reconstructed depth image. The edges extracted from rgb images match the geometry of the environments.

1) Sampling Depth Edges: Fig. 22 presents an example of reconstruction based solely on samples along the depth edges, which correspond to non-trivial pixels in the second-order derivative of the depth images. We compare the reconstructed depth images with the raw data from the ZED stereo camera, and show the statistics in Fig. 23. From Fig. 23(a) we observe that compared with the `naive` linear interpolation, `L1diag` achieves almost half the reconstruction error. In addition, the error is very small, implying almost exact recovery from samples along depth edges. On the other hand, Fig. 23(b) shows that data rate saving, defined as $1 -$ compressed data rate/uncompressed data rate, is around 70% - 85%. This implies that the same bandwidth is now able to accommodate 3x - 6x more communication channels.

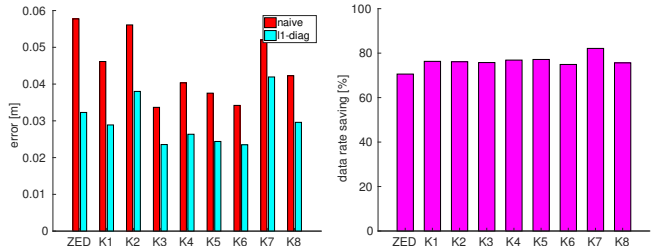

(a) (b)

Fig. 23. When only depth data along edges (extracted from depth images) are transmitted, the loss of accuracy is negligible and the data rate saving is around 80%. (a) reconstruction error. (b) data rate saving.

Note that the depth image from stereo cameras suffers from distortion and is error-prone at regions with low number of features, resulting in many edges that should not have existed, as in Fig. 22(b). For example, unnecessary edges exist on the blue floor as well as on the white walls. These unwanted edges will be treated as "important" information and thus will be sent along the communication channel, preventing further improvement on the compression rate.

2) Sampling RGB Edges: In this section, we discuss the advantages of edge extraction from the RGB images over edge extraction from the depth images. In most scenarios, depth discontinuities are reflected in appearance discontinuities. This implies that edges extracted over the RGB images should be a superset of the real edges on the depth images. By extracting edges from RGB images, we want to avoid unnecessary and erroneous edges as seen in Fig. 22(b), and thus leads to improved accuracy.

(a) RGB image (b) RGB edges (c) reconstruction

Fig. 24. An example of depth image compression with data only along RGB images: (a) RBG image, (b) edges in the RGB image, (c) `L1diag` reconstructed depth image. Undesirable edges appear due to distortion of stereo disparity images.

Fig. 24 shows the same example as in Fig. 22, but with edges extracted from the RGB images. More specifically, Fig. 24(b) shows the depth measurements along RGB edges extracted from Fig. 24(a), and Fig. 24(c) is the reconstruction result. We observe a smaller and cleaner set of edges.

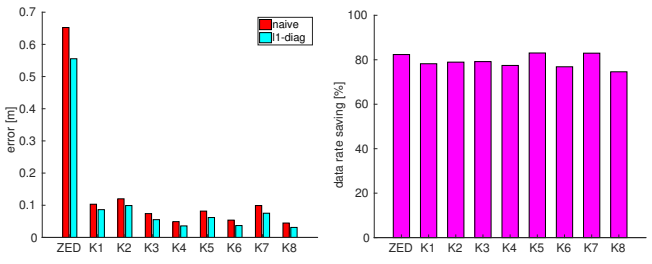

(a) (b)

Fig. 25. When only depth data along RGB edges are transmitted, the reconstruction error is minimal except for the ZED dataset. This is because of the fact that disparity images provided by the ZED stereo camera are noisy and have undesirable edges, due to inherent depth distortion of stereo-vision. In other words, edges extracted from the depth images and from the RGB images are not consistent. (a) reconstruction error. (b) data rate saving.

The reconstruction error and data rate saving statistics are also shown in Fig. 25. Note that `L1diag` still consistently outperforms the baseline `naive` on every single datasets, although they both perform poorly on the ZED stereo datasets. As discussed before, the disparity images from the ZED stereo camera suffers from distortion and does not reflect

the true geometry of the environment (see Fig. 22(b) for example). In other words, the ground truth depth is highly noisy. Therefore, even if our reconstructed depth is better than the ground truth signal, we would not expect a decline in the error metric, which is computed as the average difference between reconstruction and the ground truth.

G. Super-dense 3D reconstruction and Super-resolution Depth Imaging

In this section, we demonstrate that our algorithm can also be applied to super-resolution of depth images. This is fundamentally the same as viewing the low-resolution input depth image as measurements sampled on a grid from a higher-resolution "ground truth" and do reconstruction based on such measurements. An example has been shown in Fig. 26 with Kinect data. Fig. 26(a) is a low-resolution, 72-by-103 depth image with missing pixels (due to Kinect sensor noise). Fig. 26(b) shows the reconstructed, or in other words up-scaled, depth image from (a). The size of the reconstructed depth image is 359-by-512, and thus with an up-scale factor of 24.79. In other words, roughly 1 depth pixel in the input image translates to a 5-by-5 image patch in the up-scaled depth image. Note that all missing pixels are smoothed out, but there is also loss of information, especially small details (e.g. legs of the chairs are gone).

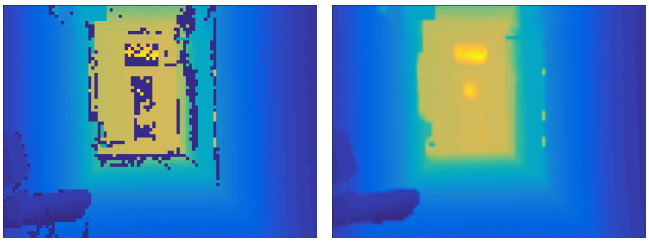

(a) 72-by-103 depth image (b) up-scaled 359-by-512 image

Fig. 26. Super-resolution of depth images. The up-scale factor is 24.79 in this example.

VIII. Conclusion and Future Work

In this work, we proposed a new approach to recover dense 2D and 3D depth profiles from sparse and incomplete measurements. We formulated depth reconstruction as a problem of finding the sparsest solution in the second-order derivative. The problem itself is NP-hard, but we used convex relaxation and reformulated the problem as constrained ℓ_1-minimization. We provided theoretical conditions under which the ground truth signal can be recovered, in both the noiseless case and when the measurements are corrupted by bounded noise.

The implementation of the presented approach is simple and straightforward, since they require solving only a convex minimization problem. In the noiseless case, the problem can be further reduced to linear programming. To accelerate the algorithm, we adapted NESTA, a first-order method for sparse recovery, to our ℓ_1-minimization. The solver significantly reduces the computational time by solving a smoothed, approximate version of the original problem.

We conducted extensive experiments on both synthetic and real data. The experimental results show that our algorithms are able to reconstruct a dense depth image from extremely low number of measurements (e.g., 40 measurements for a 100-by-100 depth image), and are robust to measurement noise in practice. The capability of properly modeling measurement noise enables a performance boost with respect to interpolation-based naive approaches. We demonstrated the power of the proposed approach in many applications, including 2D mapping, single-frame and multi-frame depth reconstruction from sparse measurements, depth image compression and decompression, as well as depth image super-resolution.

In the future, we plan to further accelerate our algorithms by parallelization on general-purpose GPUs. We would also like to apply our proposed algorithm to depth image compression in bandwidth-limited, multi-robot systems. In addition, we are interested in developing motion planning algorithms that take into account the sampling strategy of sparse depth sensors. With more intelligent motion policies, the quality of reconstructed depth images can be further improved.

References

[1] Meet piccolissimo: The worldâĂŹs smallest self-powered controllable flying vehicle. University of Pennsylvania, 2016.

[2] Pd-100 black hornet prs. proxdynamics, 2016.

[3] S.M. LaValle B. Tovar, L. Freda. Mapping and navigation from permutations of landmarks. *AMS Contemporary Mathematics*, 438:33–45, 2007.

[4] Stephen Becker, JÃľrÃťme Bobin, and Emmanuel J. Candés. Nesta: A fast and accurate first-order method for sparse recovery. *SIAM J. Imaging Sciences*, 4(1):1–39, 2011.

[5] D.P. Bertsekas, A. Nedic, and A.E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.

[6] D.P. Bertsekas and J.N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[7] P. Boning and S. Dubowsky. A study of minimal sensor topologies for space robots. *Advances in Robot Kinematics, Springer*, pages 413–422, 2006.

[8] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[9] S. Boyd and L. Vandenberghe. Subgradients. Notes for EE364b, 2006.

[10] E. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, 2006.

[11] E. Candès and M.B. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, 2008.

[12] Emmanuel Candes and Terence Tao. The dantzig selector: Statistical estimation when p is much larger than n. *The Annals of Statistics*, pages 2313–2351, 2007.

[13] Xiangyu Chen, Moju Zhao, Lingzhu Xiang, Fumihito Sugai, Hiroaki Yaguchi, Kei Okada, and Masayuki Inaba. Development of a low-cost ultra-tiny line laser range sensor. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 111–116. IEEE, 2016.

[14] C.M. da Fonseca and J. Petronilho. Explicit inverses of some tridiagonal matrices. *Linear Algebra and its Applications*, 325:7–12, 2001.

[15] Amaury Dame, Victor A Prisacariu, Carl Y Ren, and Ian Reid. Dense reconstruction using 3d object shape priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1288–1295, 2013.

[16] GCHE De Croon, C De Wagter, BDW Remes, and R Ruijsink. Sub-sampling: real-time vision for micro air vehicles. *Robotics and Autonomous Systems*, 60(2):167–181, 2012.

[17] Christophe De Wagter, Sjoerd Tijmons, Bart DW Remes, and Guido CHE de Croon. Autonomous flight of a 20-gram flapping wing mav with a 4-gram onboard stereo vision system. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4982–4987. IEEE, 2014.

[18] J. Derenick, A. Speranzon, and R. Ghrist. Homological sensing for mobile robot localization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 572–579, 2013.

[19] D. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

[20] Z. Du and G. Geng. 3-D geometric signal compression method based on compressed sensing. In *Image Analysis and Signal Processing*, pages 62–66, 2011.

[21] Marco F Duarte, Mark A Davenport, Dharmpal Takhar, Jason N Laska, Ting Sun, Kevin E Kelly, Richard G Baraniuk, et al. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):83, 2008.

[22] L.H. Erickson, J. Knuth, J.M O'Kane, and S.M. LaValle. Probabilistic localization with a blind robot. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1821–1827, 2008.

[23] Virginia Estellers, Stefano Soatto, and Xavier Bresson. Adaptive regularization with the structure tensor. *IEEE Transactions on Image Processing*, 24(6):1777–1790, 2015.

[24] Olivier D Faugeras and Francis Lustman. Motion and structure from motion in a piecewise planar environment. *International Journal of Pattern Recognition and Artificial Intelligence*, 2(03):485–508, 1988.

[25] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Birkhäuser, Boston, 2013.

[26] Brian Gerkey, Richard T Vaughan, and Andrew Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proceedings of the 11th international conference on advanced robotics*, volume 1, pages 317–323, 2003.

[27] K.J. Goldberg. Orienting polygonal parts without sensors. *Algorithmica*, 10:201–225, 1993.

[28] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming.

[29] G. Grisetti, C. Stachniss, and W. Burgard. Improved techniques for grid mapping with Rao-Blackwellized particle filters. *IEEE Trans. Robotics*, 23(1):34–46, 2007.

[30] Duncan W Haldane, Mark Plecnik, Justin K Yim, and Ronald S Fearing. Robotic vertical jumping agility via series-elastic power modulation. *Science Robotics*, 2016.

[31] W. Hong, A.Y. Yang, K. Huang, and Y. Ma. On symmetry and multiple view geometry: Structure, pose and calibration from a single image. *Intl. J. of Computer Vision*, (3), 2004.

[32] J.J Im, A. Leonessa, and A. Kurdila. A real-time data compression and occupancy grid map generation for ground-based 3d lidar data using wavelets. In *Proceedings of the ASME 2010 Dynamic Systems and Control Conference*, 2010.

[33] M. Kabanava and H. Rauhut. Cosparsity in compressed sensing. In H. Boche, R. Calderbank, G. Kutyniok, and J. Vybiral, editors, *Compressed Sensing and Its Applications*, pages 315–339. Springer, 2015.

[34] T. Kanade. Recovery of the 3-dimensional shape of an object from a single view. *Artificial Intelligence*, 17, 1981.

[35] Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 3, pages 2149–2154. IEEE, 2004.

[36] LeddarVu. Leddarvu solid-state lidar platform. http://leddartech.com/, 2016.

[37] X. Liang, X. Ren, Z. Zhang, and Y. Ma. Repairing sparse low-rank texture. In *European Conf. on Computer Vision (ECCV)*, pages 482–495, 2012.

[38] Lee-Kang Liu, Stanley H Chan, and Truong Q Nguyen. Depth reconstruction from sparse samples: Representation, algorithm, and sampling. *IEEE Transactions on Image Processing*, 24(6):1983–1996, 2015.

[39] Jiajun Lu and David Forsyth. Sparse depth super resolution. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2245–2253. IEEE, 2015.

[40] Fangchang Ma, Luca Carlone, Ulas Ayaz, and Sertac Karaman. Sparse sensing for resource-constrained depth reconstruction. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 96–103. IEEE, 2016.

[41] D. Marinakis and G. Dudek. Occam's razor applied to network topology inference. *IEEE Trans. Robotics*, 24(2):293–306, April 2008.

[42] D. Marinakis, P. Giguere, and G. Dudek. Learning network topology from simple sensor data. In *Prof. Canadian Conference on Artificial Intelligence*, 2002.

[43] M. Milford. Vision-based place recognition: how low can you go? *Intl. J. of Robotics Research*, 32:766–789, 2013.

[44] R. Mur-Artal, J.M.M. Montiel, and J.D. Tardós. ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE Trans. Robotics*, 31(5):1147–1163, 2015.

[45] S. Nam, M.E. Davies, M. Elad, and R. Gribonval. The cosparse analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2013.

[46] D. Needell and R. Ward. Stable image reconstruction using total variation minimization. *SIAM J. Imaging Sci.*, 6(2):1035–1058, 2013.

[47] E. Nelson and N. Michael. Information-theoretic occupancy grid compression for high-speed information-based exploration. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 4976–4982, 2015.

[48] Yu Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, May 2005.

[49] Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence o (1/k2). In *Doklady an SSSR*, volume 269, pages 543–547, 1983.

[50] R.A. Newcombe, S.J. Lovegrove, and A.J. Davison. DTAM: Dense tracking and mapping in real-time. In *Intl. Conf. on Computer Vision (ICCV)*, pages 2320–2327, Barcelona, Spain, Nov 2011.

[51] J. M. O'Kane and S. M. LaValle. On comparing the power of mobile robots. In *Robotics: Science and Systems (RSS)*, 2006.

[52] J. M. O'Kane and S. M. LaValle. Localization with limited sensing. *IEEE Trans. Robotics*, 23:704–716, 2007.

[53] S. OâĂŹCallaghan and F. Ramos. Gaussian process occupancy maps for dynamic environments. In *14th International Symposium on Experimental Robotics*, 2014.

[54] David B Phillips, Ming-Jie Sun, Jonathan M Taylor, Matthew P Edgar, Stephen M Barnett, Graham G Gibson, and Miles J Padgett. Adaptive foveated single-pixel imaging with dynamic super-sampling. *arXiv preprint arXiv:1607.08236*, 2016.

[55] S. Pillai, S. Ramalingam, and J. Leonard. High-performance and tunable stereo reconstruction. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2016.

[56] P. Piniés, L.M. Paz, and P. Newman. Too much TV is bad: dense reconstruction from sparse laser with non-convex regularisation. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015.

[57] Christopher Poulton and Michael Watts. Mit and darpa pack lidar sensor onto single chip. IEEE Spectrum, 2016.

[58] F. Ramos and L. Ott. Hilbert maps: scalable continuous occupacy mapping with stochastic gradient descent. In *Robotics: Science and Systems (RSS)*, 2015.

[59] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, August 2010.

[60] D. Reddy, A. C. Sankaranarayanan, V. Cevher, and R. Chellappa. Compressed sensing for multi-view tracking and 3-D voxel reconstruction. In *Intl. Conf. on Image Processing (ICIP)*, pages 221–224, 2008.

[61] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268, November 1992.

[62] M. Ruhnke, L. Bo, D. Fox, and W. Burgard. Compact rgbd surface models based on sparse coding. In *Nat. Conf. on Artificial Intelligence (AAAI)*, 2013.

[63] M. Ruhnke, L. Bo, D. Fox, and W. Burgard. Hierarchical sparse coded surface models. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 6238–6243, 2014.

[64] N. Shrivastava, R.M. Madhow, and S. Suri. Target tracking with binary proximity sensors: fundamental limits, minimal descriptions, and algorithms. In *Proc. of the 4th Intl. Conf. on Embedded Networked Sensor Systems*, pages 251–264, 2006.

[65] Baoqing Sun, Matthew P Edgar, Richard Bowman, Liberty E Vittert, Stuart Welsh, A Bowman, and MJ Padgett. 3d computational imaging with single-pixel detectors. *Science*, 340(6134):844–847, 2013.

[66] S. Suri, E. Vicari, and P.Widmayer. Simple robots with minimal sensing: From local visibility to global geometry. *Intl. J. of Robotics Research*, 27, 2008.
[67] TeraRanger. Teraranger one. http://www.teraranger.com/, 2016.
[68] B. Tovar, F. Cohen, L. Bobadilla, J. Czarnowski, and S. M. LaValle. Combinatorial filters: Sensor beams, obstacles, and possible paths. *ACM Transactions on Sensor Networks*, 10(3), 2014.
[69] B. Tovar, F. Cohen, and S. M. LaValle. Sensor beams, obstacles, and possible paths. In *Intl. Workshop on the Algorithmic Foundations of Robotics*, 2008.
[70] B. Tovar, L. Freda, and S. M. LaValle. Learning combinatorial map information from permutations of landmarks. *Intl. J. of Robotics Research*, 30(9):1143–1156, 2011.
[71] B. Tovar and S. M. LaValle. Searching and mapping among indistinguishable convex obstacles. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010.
[72] B.E. Usevitch. A tutorial on modern lossy wavelet image compression: Foundations of jpeg 2000. *Signal Processing Magazine*, 18(5):22–35, 2001.
[73] S. Vaiter, G. Peyré, Ch. Dossal, and J. Fadili. Robust sparse analysis regularization. *IEEE Transactions on information theory*, 59(4):2001–2016, 2013.
[74] T. Whelan, S. Leutenegger, R. Salas-Moreno, B. Glocker, and A. Davison. ElasticFusion: Dense SLAM without a pose graph. In *Robotics: Science and Systems (RSS)*, 2015.
[75] R. Wood. RoboBees project. Harvard University, 2015.

APPENDIX

A. Some Useful Lemmas

We introduce some technical lemmas which simplify the derivations in the following sections.

Lemma 26 (Null space of A)**.** *Consider the sparse sampling matrix $A \in \mathbb{R}^{m\times n}$. The null space of A is spanned by the rows of the matrix $N \doteq \mathbf{I}_{\overline{\mathcal{M}}} \in \mathbb{R}^{\bar{m}\times n}$ (with $\bar{m} \doteq n-m$). Moreover, the action of the matrix N on a vector v and on a matrix V of suitable dimensions is such that $Nv = v_{\overline{\mathcal{M}}}$ and $NV = V_{\overline{\mathcal{M}}}$.*

Proof. Denote the i-th standard basis vector as e_i. Each row of $A \in \mathbb{R}^{m\times n}$ is equal to e_i^T for some $i \in \mathcal{M}$, hence A has rank m. Since the sets $\mathcal{M}$ and $\overline{\mathcal{M}}$ are disjoint and are such that $\mathcal{M} \cup \overline{\mathcal{M}} = \{1, \dots, n\}$, it follows that $AN^\mathsf{T} = 0$ (entries of AN^T have the form $e_i^\mathsf{T} e_j$ which is zero for $\mathcal{M} \ni i \neq j \in \overline{\mathcal{M}}$) and N^T has rank $\bar{m} = n - m$. This proves that the rows of N span the null space of A. Since each row of N is e_j^T for some $j \in \overline{\mathcal{M}}$, the claims $Nv = v_{\overline{\mathcal{M}}}$ and $NV = V_{\overline{\mathcal{M}}}$ easily follow. □

Lemma 27 (Symmetric Tridiagonal Toeplitz matrix)**.** *Let T denote a Toeplitz symmetric tridiagonal Toeplitz matrix with diagonal entries equal to -2 and off-diagonal entries equal to 1:*

$$T \doteq \begin{bmatrix} -2 & 1 & 0 & 0 & \dots & 0 \\ 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & 1 & -2 & -1 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 0 & 1 & -2 \end{bmatrix} \tag{23}$$

Then the following claims hold:

(i) T is invertible;
(ii) all the entries in the first and in the last column of T^{-1} are negative and have absolute value smaller than 1;
(iii) let $v \in \mathbb{R}^d$ be defined as $v \doteq [1\ 0\ \dots\ 0\ 1]^\mathsf{T}$, then $T^{-1}v = -\mathbf{1}$.

Proof. Invertibility follows from [14, Corollary 4.2], which also reports the explicit form of the inverse of a Toeplitz matrix. We report the inverse here, tailoring it to our matrix. The entry in row i and column j of T^{-1} is:

$$(T^{-1})_{ij} = \begin{cases} (-1)^{2i-1}\frac{i(n-j+1)}{n+1} & \text{if } i \leq j \\ (-1)^{2j-1}\frac{j(n-i+1)}{n+1} & \text{if } i > j \end{cases} \tag{24}$$

By inspection one can see that the first column ($j = 1$) and the last column $j = d$ are all negative and have absolute value smaller than 1. The last claim can be proved by observing that $T\mathbf{1} = -v$ and the matrix is invertible. □

Lemma 28 (Null Space of D)**.** *Given a 2nd-order difference operator $D \in \mathbb{R}^{(n-2)\times n}$, defined as in* (5)*, the null space of D is spanned by the following vectors:*

$$v_1 = \mathbf{1}_n \qquad v_2 = [1\ 2\ \dots\ n]^\mathsf{T} \tag{25}$$

Proof. By inspection one can see that $Dv_1 = Dv_2 = \mathbf{0}$. Moreover, the rank of D is $n-2$ and v_1 and v_2 are two linearly independent vectors, which proves the claim. □

B. Equivalence of Constrained and Unconstrained Forms

Here we show that Problem (L1$_D^\varepsilon$) is equivalent to:

$$\min_z \frac{1}{2}\|Az-y\|_2^2+\lambda\|Dz\|_1 \qquad \text{(IP)}$$

for a suitable choice of the parameter $\lambda > 0$; the problems are equivalent in the sense that they admit the same solution $z^\star$. We first rewrite (L1$_D^\varepsilon$) as:

$$\min_z \|Dz\|_1 \quad \text{subject to} \quad \|Az-y\|_2^2 \leq \varepsilon^2 \qquad (26)$$

where we noted that squaring both (nonnegative) terms in the inequality does not alter the feasible set. We prove the equivalence as follows. We compute the Lagrangian of (L1$_D^\varepsilon$) and we show that it reduces to the objective in (IP); moreover, since (L1$_D^\varepsilon$) is convex and admits a strictly feasible solution, strong duality holds, and the optimal solution of (L1$_D^\varepsilon$) is a saddle point of the Lagrangian, which in turns is a minimizer of (IP) for a specific λ, which is connected to the solution of the dual problem of (L1$_D^\varepsilon$).

Lagrangian and equivalence. We start by rewriting (26) in a form that is more amenable to derive the dual problem:

$$\begin{array}{ll} \min_z & \|a\|_1 \\ \text{subject to} & \|b\|_2^2 \leq \varepsilon^2 \\ & Dz = a \\ & Az - y = b \end{array} \qquad (27)$$

where we included slack variables $a \in \mathbb{R}^{n-2}$ and $b \in \mathbb{R}^m$. The Lagrangian of problem (27) is:

$$\begin{aligned} \mathcal{L}(z,a,b,\eta_a,\eta_b,\eta_\varepsilon) = \|a\|_1+\eta_\varepsilon(\|b\|_2^2-\varepsilon^2) + \\ +\eta_a^\mathsf{T}(Dz-a)+\eta_b^\mathsf{T}(Az-y-b) \end{aligned} \qquad (28)$$

where $\eta_\varepsilon \in \mathbb{R}$ (and $\eta_\varepsilon \geq 0$), $\eta_a \in \mathbb{R}^{n-2}$, and $\eta_b \in \mathbb{R}^m$ are the *dual variables* or *Lagrange multipliers*. The dual problem is defined as:

$$\max_{\eta_a,\eta_b,\eta_\varepsilon} \inf_{z,a,b} \mathcal{L}(z,a,b,\eta_a,\eta_b,\eta_\varepsilon) \quad \text{subject to} \quad \eta_\varepsilon \geq 0 \qquad \text{(dualBP)}$$

Let us call $\eta_\varepsilon^\star, \eta_a^\star, \eta_b^\star$ a solution of the dual problem (dualBP) and let $z^\star, a^\star, b^\star$ be a solution of (L1$_D^\varepsilon$). Now, we know that (L1$_D^\varepsilon$) is convex and admits a strictly feasible solution, hence strong duality attains. Therefore, $z^\star, a^\star, b^\star$ must be in the set of minimizers of $\mathcal{L}(z,a,b,\eta_\varepsilon^\star,\eta_a^\star,\eta_b^\star)$ [8, Page 248], i.e., we can compute $z^\star, a^\star, b^\star$ by solving the following:

$$\begin{aligned} \min_{z,a,b} \mathcal{L}(z,a,b,\eta_\varepsilon^\star,\eta_a^\star,\eta_b^\star) = \min_{z,a,b} \|a\|_1 + \\ +\eta_\varepsilon^\star(\|b\|_2^2-\varepsilon^2)+\eta_a^{\star\mathsf{T}}(Dz-a)+\eta_b^{\star\mathsf{T}}(Az-y-b) \end{aligned} \qquad (29)$$

Since $z^\star, a^\star, b^\star$ must be feasible for (27), eq. (29) simplifies to:

$$\min_z \|Dz\|_1+\eta_\varepsilon^\star\|Az-y\|_2^2 \qquad (30)$$

where we also removed the constant term $\eta_\varepsilon^\star\varepsilon^2$, which is inconsequential. Now we can see that (30) is the same as (IP) for $\lambda = 1/2\eta_\varepsilon^\star$. This proves that $z^\star$ can be computed by solving (IP), which proves our claim.

Dual optimal solution. For completeness we conclude this section by providing a more explicit form for the dual problem (dualBP). We start by writing (dualBP) explicitly:

$$\max_{\eta_a,\eta_b,\eta_\varepsilon} \inf_{z,a,b} \|a\|_1+\eta_\varepsilon(\|b\|_2^2-\varepsilon^2)+\eta_a^\mathsf{T}(Dz-a)+\eta_b^\mathsf{T}(Az-y-b) \quad \text{subject to} \quad \eta_\varepsilon \geq 0$$

We rearrange the terms in the objective as follows:

$$\begin{aligned} \max_{\eta_a,\eta_b,\eta_\varepsilon} \inf_a (\|a\|_1-\eta_a^\mathsf{T}a)+\inf_b \left(\eta_\varepsilon\|b\|_2^2-\eta_b^\mathsf{T}b\right)+ \\ \inf_z (\eta_a^\mathsf{T}Dz+\eta_b^\mathsf{T}Az)-\eta_\varepsilon\varepsilon^2-\eta_b^\mathsf{T}y = \\ -\sup_a (\eta_a^\mathsf{T}a-\|a\|_1)-2\eta_\varepsilon\sup_b \left(\frac{\eta_b^\mathsf{T}}{2\eta_\varepsilon}b-\frac{1}{2}\|b\|_2^2\right)+ \\ \inf_z (\eta_a^\mathsf{T}D+\eta_b^\mathsf{T}A)z-\eta_\varepsilon\varepsilon^2-\eta_b^\mathsf{T}y \end{aligned} \qquad (31)$$

Now we note that the infimum in z is:

$$\inf_z (\eta_a^\mathsf{T}D+\eta_b^\mathsf{T}A)z = \begin{cases} 0 & \text{if } \eta_a^\mathsf{T}D+\eta_b^\mathsf{T}A = 0 \\ -\infty & \text{otherwise} \end{cases} \qquad (32)$$

Moreover, $\sup_a (\eta_a^\mathsf{T}a-\|a\|_1)$ is, by definition, the conjugate function of the ℓ_1 norm, hence it holds [8, Example 3.26]:

$$-\sup_a (\eta_a^\mathsf{T}a-\|a\|_1) = \begin{cases} 0 & \text{if } \|\eta_a\|_\infty \leq 1 \\ -\infty & \text{otherwise} \end{cases} \qquad (33)$$

Furthermore, $\sup_b \left(\frac{\eta_b^\mathsf{T}}{2\eta_\varepsilon}b-\frac{1}{2}\|b\|_2^2\right)$ is the conjugate function of $\frac{1}{2}\|\cdot\|_2^2$, hence it holds [8, Example 3.27]:

$$-\sup_b \left(\frac{\eta_b^\mathsf{T}}{2\eta_\varepsilon}b-\frac{1}{2}\|b\|_2^2\right) = -\frac{1}{2}\left\|\frac{\eta_b^\mathsf{T}}{2\eta_\varepsilon}\right\|_2^2 \overset{\eta_\varepsilon\geq 0}{=} -\frac{\|\eta_b\|_2^2}{8\eta_\varepsilon^2} \qquad (34)$$

Since we look for a maximum, we can safely rule out the cases in which the cost drifts to minus infinity. Therefore, using (32)-(34), the dual problem simplifies to:

$$\begin{array}{ll} \max\limits_{\eta_a,\eta_b,\eta_\varepsilon} & -\dfrac{\|\eta_b\|_2^2}{4\eta_\varepsilon}-\eta_b^\mathsf{T}y-\eta_\varepsilon\varepsilon^2 \\ \text{subject to} & \eta_\varepsilon \geq 0 \\ & \eta_a^\mathsf{T}D+\eta_b^\mathsf{T}A = 0 \\ & \|\eta_a\|_\infty \leq 1 \end{array} \qquad (35)$$

For any choice of η_a and η_b, the minimizer of η_ε is the one making the gradient of the cost equal to zero, i.e.,

$$\nabla_{\eta_\varepsilon}\left(-\frac{\|\eta_b\|_2^2}{4\eta_\varepsilon^\star}-\eta_b^\mathsf{T}y-\eta_\varepsilon^\star\varepsilon^2\right) = \frac{\|\eta_b\|_2^2}{4{\eta_\varepsilon^\star}^2}-\varepsilon^2 = 0 \qquad (36)$$

which implies $\eta_\varepsilon^\star = \pm\frac{\|\eta_b\|_2}{2\varepsilon}$. Since we need $\eta_\varepsilon \geq 0$, it follows

$$\eta_\varepsilon^\star = \frac{\|\eta_b\|_2}{2\varepsilon} \qquad (37)$$

Substituting $\eta_\varepsilon^\star$, which is the optimal η_ε for any choice of η_a and η_b, problem (35) becomes:

$$\begin{array}{ll} \max\limits_{\eta_a,\eta_b} & -\varepsilon\|\eta_b\|_2-\eta_b^\mathsf{T}y \\ \text{subject to} & \eta_a^\mathsf{T}D+\eta_b^\mathsf{T}A = 0 \\ & \|\eta_a\|_\infty \leq 1 \end{array} \qquad (38)$$

which is a convex program whose solution is $(\eta_a^\star, \eta_b^\star)$.

C. Proof of Proposition 7

In this section we prove that in 2D depth estimation, if we sample only the corners of the profile (and the first and the last entry), then $C_{\text{er}} = 1$. Moreover, we prove that if we sample the corners and their neighbors we have $C_{\text{er}} = 0$. We start by rewriting (12) in a more convenient form. Using Lemma 26, we know that $N(D_{\mathcal{J}})^{\mathsf{T}} = [D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}}$. In words, $D_{\mathcal{J}}$ selects the rows of D at indices $\mathcal{J}$, or equivalently the columns of D^{T}. Similarly the multiplication by N selects the rows of D^{T} at indices $\overline{\mathcal{M}}$. Similarly, $N(D_{\mathcal{I}})^{\mathsf{T}} = [D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}$. Using these relations, (12) simplifies to:

$$C_{\text{er}} = \|([D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}})^{\dagger}[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}\|_{\infty\to\infty} < 1 \tag{39}$$

Since $\mathcal{I} \cup \mathcal{J} = \{1, \dots, n\}$, it is clear that $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}}$ and $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}$ are disjoint sets of columns of the matrix $[D^{\mathsf{T}}]_{\overline{\mathcal{M}}}$.

Let us start with the first claim: $C_{\text{er}} = 1$ whenever we sample the corners, the first, and the last point. We will make an extensive use of the structure of the matrix D^{T} which is the transpose of (5). To give a more intuitive understanding of the proof we provide a small example of D^{T} with $n = 12$:

$D^{\mathsf{T}} =$

1	0	0	0	0	0	0	0	0	0
−2	1	0	0	0	0	0	0	0	0
1	−2	1	0	0	0	0	0	0	0
0	1	−2	1	0	0	0	0	0	0
0	0	1	−2	1	0	0	0	0	0
0	0	0	1	−2	1	0	0	0	0
0	0	0	0	1	−2	1	0	0	0
0	0	0	0	0	1	−2	1	0	0
0	0	0	0	0	0	1	−2	1	0
0	0	0	0	0	0	0	1	−2	1
0	0	0	0	0	0	0	0	1	−2
0	0	0	0	0	0	0	0	0	1

$D_{\mathcal{I}}^{\mathsf{T}}$

(40)

The matrix $[D^{\mathsf{T}}]_{\overline{\mathcal{M}}}$ is obtained from D^{T} after removing the rows at indices in the sample set $\mathcal{M}$: these "deleted" rows are shown in gray in (40). In particular, according to the assumptions of the first claim of Proposition 7, $\mathcal{M}$ contains the first and the last sample (first and last row in D^{T}) plus intermediate rows corresponding to corners (two intermediate gray rows in the figure). Now we note that the matrix $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}$ selects the columns with indices in $\mathcal{I}$ from $[D^{\mathsf{T}}]_{\overline{\mathcal{M}}}$. In figure, the columns that form $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}$ are shown in dashed red boxes. The position of these columns is dictated by the position of the corners, hence if the i-th row corresponds to a corner, than column $i-1$ belongs to $\mathcal{I}$.

Three considerations are in order now. First, the matrix $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}}$ is a block-diagonal square matrix with diagonal blocks being Toeplitz matrices (*cf.* with (23)). Second, the matrix is invertible (follows from the first claim of Lemma 27). Third, the matrix $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}$ only contains 0 and 1 in suitable positions. Therefore, the matrix $([D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}})^{\dagger}[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}} = ([D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}})^{-1}[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}$ has the following block structure:

$$\begin{bmatrix} T_1^{-1}\begin{bmatrix} 0 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & \dots & 0 & 0 \end{bmatrix} \\ T_2^{-1}\begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 1 & \dots & 0 & 0 \end{bmatrix} \\ \vdots \\ T_K^{-1}\begin{bmatrix} 0 & 0 & \dots & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 \end{bmatrix} \end{bmatrix} = \begin{bmatrix} T_1^{-1}R_1 \\ T_2^{-1}R_2 \\ \vdots \\ T_K^{-1}R_K \end{bmatrix} \tag{41}$$

where $T_1, \dots, T_K$ are Toeplitz matrices of suitable dimensions and R_i contain at most two nonzero elements (equal to 1) in the first and the last row. Since for a matrix M, $\|M\|_{\infty\to\infty}$ is the maximum of the ℓ_1-norm of each row, we only need to demonstrate that the maximum ℓ_1-norm of the rows of $T_i^{-1}R_i$ is no larger than 1 for all i. The action of R_i on T_i^{-1} is to select the first and/or the last column of T_i^{-1} (depending on where the 1 appears). In particular $T_1^{-1}R_1$ is zero everywhere, except the fist column which is equal to the last column of T_1^{-1}. From Lemma 27(ii) we know that the entries in the last column have magnitude smaller than 1 hence it follows that $\|T_1^{-1}R_1\|_{\infty\to\infty} < 1$. A similar argument holds for the last block, hence $\|T_K^{-1}R_K\|_{\infty\to\infty} < 1$. For the intermediate blocks $T_i^{-1}R_i$, $1 < i < K$, it can be seen that $\|T_i^{-1}R_i\|_{\infty\to\infty} = \|T_i^{-1}v\|_{\infty}$, with $v \doteq [1\ 0\ \dots\ 0\ 1]^{\mathsf{T}}$; this follows from the fact that R_i selects the first and the last columns of T_i^{-1} which have negative entries due to Lemma 27(ii). Using Lemma 27(iii) we know that $T^{-1}v = -\mathbf{1}$, from which it follows $\|T_i^{-1}R_i\|_{\infty\to\infty} = \|T_i^{-1}v\|_{\infty} = 1$. This proves the first claim.

The proof of the second claim ($C_{\text{er}} = 0$ when we sample the corners and their neighbors) is much simpler. Sampling the neighbors corresponds to deleting the rows contiguous to each "corner" from D^{T}. From (40) the reader can easily see that this choice makes $[D^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}} = \mathbf{0}$, which in turns implies $C_{\text{er}} = 0$. □

D. Proof of Proposition 8

In this section we show that when we sample the edges and the corresponding vertical and horizontal neighbors, then:

$$C_{\text{er}} = \|(N(\Delta_{\mathcal{J}})^{\mathsf{T}})^{\dagger}N(\Delta_{\mathcal{I}})^{\mathsf{T}})\|_{\infty\to\infty} = 0 \tag{42}$$

which implies exact recovery of the original depth profile according to Proposition 6. As in C we rewrite the condition (42) as:

$$C_{\text{er}} = \|([\Delta^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{J}})^{\dagger}[\Delta^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}}\|_{\infty\to\infty} = 0 \tag{43}$$

The proof proceeds along the same line of the proof of the second claim in Proposition 7. By observing the structure of Δ^{T}, we realize that sampling the edges and the corresponding vertical and horizontal neighbors, makes $[\Delta^{\mathsf{T}}]_{\overline{\mathcal{M}},\mathcal{I}} = \mathbf{0}$, which in turns implies $C_{\text{er}} = 0$.

E. Proof of Proposition 9

In this section we establish necessary and sufficient conditions for an estimate $z^\star$ to be in the set $\mathcal{S}^\star$ of optimal solutions of problem ($\mathrm{L1}_D$). The proof is identical for the 3D case in Corollary 10 (substituting D with Δ), hence we restrict ourselves to the 2D case. We start by rewriting ($\mathrm{L1}_D$) as:

$$\min_z \|Dz\|_1 + \chi_{\{Az=y\}} =: \min_z f(z) \tag{44}$$

where $\chi_{\{Az=y\}}$ is the indicator function of the set $\{z : Az = y\}$, which is zero whenever $Az = y$ and $+\infty$ otherwise. Since $Az = y$ defines a convex (affine) set, (44) is convex. In the following we make extensive use of subgradients of convex functions. We refer the reader to [6, §4] for a comprehensive treatment and to [9] for a quick introduction.

A point $z^\star$ is a minimizer of a convex function f if and only if f is subdifferentiable at $z^\star$ and the zero vector belongs to the set of subgradients of f, i.e., $\mathbf{0} \in \partial f(z^\star)$. The set of subgradients is also called the *subdifferential*. The subdifferential of a sum of functions is the sum of the subdifferentials, therefore

$$\partial f(z^\star) = \partial(\|Dz\|_1)(z^\star) + \partial(\chi_{\{Az=y\}})(z^\star) \tag{45}$$

In the following we compute each subdifferential in (45). Let us call $\mathcal{I}$ the support set of the vector $Dz^\star$, and recall that, given a vector v, we denote with $v_\mathcal{I}$ the subvector of v including the entries of v at indices in $\mathcal{I}$. Using [9, Page 5]:

$$\partial(\|Dz\|_1)(z^\star) = \{D^\mathsf{T}u \in \mathbb{R}^n : u_\mathcal{I} = \operatorname{sign}(Dz^\star)_\mathcal{I}, \|u_\mathcal{J}\|_\infty \leq 1\} \tag{46}$$

The second subdifferential in (45) is [5, Page 254]:

$$\partial(\chi_{\{Az=y\}})(z^\star) = \{g \in \mathbb{R}^n : g^\mathsf{T}z^\star \geq g^\mathsf{T}r,\ \forall r \text{ s.t. } Ar = y\} \tag{47}$$

To get a better understanding of the set in (47), we note that every solution $r \in \mathbb{R}^n$ of the overdetermined linear system $Ar = y$ can be written as a vector that satisfies the linear system, plus a vector that is in the null space of A. Now we know that $z^\diamond$, the vector that generated the data y, satisfies $Az^\diamond = y$. Therefore, we rewrite (47) as:

$$\partial(\chi_{\{Az=y\}})(z^\star) = \{g \in \mathbb{R}^n : g^\mathsf{T}z^\star \geq g^\mathsf{T}(z^\diamond + \tilde{r}),\ \forall \tilde{r} \in \ker(A)\} \tag{48}$$

where $\ker(A)$ denotes the kernel of A. From Lemma 26 we know that the kernel A is spanned by the matrix R (defined in the lemma), hence (48) further simplifies to:

$$\partial(\chi_{\{Az=y\}})(z^\star) = \{g \in \mathbb{R}^n : g^\mathsf{T}z^\star \geq g^\mathsf{T}(z^\diamond + Rw),\ \forall w \in \mathbb{R}^{\bar{m}}\} \tag{49}$$

Rearranging the terms:

$$\partial(\chi_{\{Az=y\}})(z^\star) = \{g \in \mathbb{R}^n : g^\mathsf{T}(z^\star - z^\diamond) \geq (R^\mathsf{T}g)^\mathsf{T}w,\ \forall w \in \mathbb{R}^{\bar{m}}\} \tag{50}$$

From the second claim of Lemma 26 we know that $R^\mathsf{T}g = g_{\overline{\mathcal{M}}}$; moreover, we observe that if an element of $g_{\overline{\mathcal{M}}}$ is different from zero, then we can pick an arbitrarily large w that falsifies the inequality, therefore, it must hold $g_{\overline{\mathcal{M}}} = \mathbf{0}$. Therefore, we rewrite (50) as:

$$\partial(\chi_{\{Az=y\}})(z^\star) = \{g \in \mathbb{R}^n : g^\mathsf{T}(z^\star - z^\diamond) \geq 0,\ g_{\overline{\mathcal{M}}} = \mathbf{0}\} \tag{51}$$

Now we split the product $g^\mathsf{T}(z^\star - z^\diamond)$ as $g^\mathsf{T}_{\overline{\mathcal{M}}}(z^\star - z^\diamond)_{\overline{\mathcal{M}}} + g^\mathsf{T}_{\mathcal{M}}(z^\star - z^\diamond)_{\mathcal{M}}$ and note that $g_{\overline{\mathcal{M}}} = 0$. Moreover, for any feasible $z^\star$, the i-th entry of $z^\star - z^\diamond$ is zero for all $i \in \mathcal{M}$, which implies $g^\mathsf{T}_{\mathcal{M}}(z^\star - z^\diamond)_{\mathcal{M}} = \mathbf{0}$. Therefore, the inequality $g^\mathsf{T}(z^\star - z^\diamond) \geq 0$ vanishes and we remain with:

$$\partial(\chi_{\{Az=y\}})(z^\star) = \{g \in \mathbb{R}^n : g_{\overline{\mathcal{M}}} = \mathbf{0}\} \tag{52}$$

Substituting (52) and (46) back into (45), we obtain:

$$\partial f(z^\star) = \{D^\mathsf{T}u + g \ : \ u_\mathcal{I} = \operatorname{sign}(Dz^\star)_\mathcal{I},\ \|u_\mathcal{J}\|_\infty \leq 1,\ g_{\overline{\mathcal{M}}} = \mathbf{0}\} \tag{53}$$

We can now use the subdifferential (53) to describe the optimal solution set $\mathcal{S}^\star$ of ($\mathrm{L1}_D$); as mentioned earlier in this section, $z^\star$ is optimal if and only if zero is a subgradient, therefore $\mathcal{S}^\star$ is defined as:

$$\begin{aligned}\mathcal{S}^\star = \{z^\star \ : \ & \exists u \in \mathbb{R}^{n-2}, g \in \mathbb{R}^n, \text{ such that} \\ & D^\mathsf{T}u + g = \mathbf{0}, \\ & u_\mathcal{I} = \operatorname{sign}(Dz^\star)_\mathcal{I},\ \ \|u_\mathcal{J}\|_\infty \leq 1, \\ & g_{\overline{\mathcal{M}}} = \mathbf{0}\}\end{aligned} \tag{54}$$

We note that the constraints $D^\mathsf{T}u + g = \mathbf{0}$ and $g_{\overline{\mathcal{M}}} = \mathbf{0}$ can be written compactly as $[D^\mathsf{T}u]_{\overline{\mathcal{M}}} = \mathbf{0}$, which is the same as $(D^\mathsf{T})_{\overline{\mathcal{M}}}\, u = \mathbf{0}$. This allows rewriting (54) as:

$$\begin{aligned}\mathcal{S}^\star = \{z^\star \ : \ & \exists u \in \mathbb{R}^{n-2}, \text{ such that} \\ & (D^\mathsf{T})_{\overline{\mathcal{M}}}\, u = \mathbf{0}, \\ & u_\mathcal{I} = \operatorname{sign}(Dz^\star)_\mathcal{I},\ \ \|u_\mathcal{J}\|_\infty \leq 1\}\end{aligned} \tag{55}$$

which coincides with the optimality condition of Proposition 9, proving the claim. □

F. Proof of Theorem 13

Theorem 13 says that sign consistency of z is a necessary and sufficient condition for z to be in the solution set of ($\mathrm{L1}_D$). In the following we denote with $\mathcal{SC}$ the set of sign consistent signals which are feasible for ($\mathrm{L1}_D$). Moreover, we denote with $\mathcal{S}^\star$ the set of optimal solutions of ($\mathrm{L1}_D$). The proof relies on the optimality conditions of Proposition 9, which we recall here: a signal z is in the solution set $\mathcal{S}^\star$ if and only if there exists a $u \in \mathbb{R}^{n-2}$ such that

$$(D^\mathsf{T})_{\overline{\mathcal{M}}}\, u = 0 \ \text{ and } \ u_\mathcal{I} = \operatorname{sign}(Dz)_\mathcal{I} \ \text{ and } \ \|u_\mathcal{J}\|_\infty \leq 1 \tag{56}$$

where $\mathcal{I}$ is the support set of the vector Dz. Before proving that $z \in \mathcal{SC} \Leftrightarrow z \in \mathcal{S}^\star$, we need a better understanding of the structure of the matrix $(D^\mathsf{T})_{\overline{\mathcal{M}}}$. We note that, when taking twin samples, the matrix $(D^\mathsf{T})_{\overline{\mathcal{M}}}$ is obtained from D^T by removing pairs of consecutive rows. For instance,

considering a problem with $n = 12$, the product $(D^\mathsf{T})_{\overline{\mathcal{M}}} u$ becomes:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -2 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -2 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ u_5 \\ u_6 \\ u_7 \\ u_8 \\ u_9 \\ u_{10} \end{bmatrix} \tag{57}$$

where gray rows are the ones we "removed" from D^T to obtain $(D^\mathsf{T})_{\overline{\mathcal{M}}}$. By observing (57), the reader can verify that the resulting matrix $(D^\mathsf{T})_{\overline{\mathcal{M}}}$ is block diagonal (in general will have more than 2 diagonal blocks), and each block is a 2nd-order difference operator like (5) of suitable size. We denote the diagonal blocks as $D^{(1)}, D^{(2)}, \dots, D^{(K)}$ ($K = 2$ in the example of eq. (57)). This also induces a partition in the vector u, which can be split vertically as $u = [u_{S_1}\ u_{S_2}\ \dots\ u_{S_K}]$. Geometrically, the block diagonal structure that arises means that each segment between consecutive twin samples can be studied independently. Therefore, the condition (56) can be written as:

$$D^{(k)}\, u_{S_k} = 0 \quad \text{and} \quad u_{\mathcal{I}_k} = \text{sign}(Dz)_{\mathcal{I}_k} \quad \text{and} \quad \|u_{S_k}\|_\infty \leq 1 \tag{58}$$

for $k = 1, \dots, K$, where $\mathcal{I}_k \subseteq S_k$ are entries in the support set $\mathcal{I}$ that fall within the set S_k. With this machinery we are ready to prove Theorem 13.

Let us start with the implication $z \in \mathcal{SC} \Rightarrow z \in \mathcal{S}^\star$. We first consider the case in which all signs are consistent, i.e., Definition 11(i). This means that within each set S_k, $\text{sign}(Dz)_{\mathcal{I}_k} = +\mathbf{1}$ or $\text{sign}(Dz)_{\mathcal{I}_k} = -\mathbf{1}$. Without loss of generality, assume $\text{sign}(Dz)_{\mathcal{I}_k} = +\mathbf{1}$. Then, we can see that selecting $u_{S_k} = +\mathbf{1}$ satisfies $u_{\mathcal{I}_k} = \text{sign}(Dz)_{\mathcal{I}_k}$ and $\|u_{S_k}\|_\infty \leq 1$. Moreover, since $\mathbf{1}$ is in the null space of $D^{(k)}$ (see Lemma 28), it follows $D^{(k)} u_{S_k} = 0$, proving the claim. To complete the demonstration of $z \in \mathcal{SC} \Rightarrow z \in \mathcal{S}^\star$ we consider the case in which there is there is a sign change at the boundary of each segment, while all signs are zero in the interior (Definition 11(ii)). In this case the condition $u_{\mathcal{I}_k} = \text{sign}(Dz)_{\mathcal{I}_k}$ imposes that the first and the last elements of $u_{\mathcal{I}_k}$ are $+1$ and -1 (or -1 and $+1$) respectively. Without loss of generality, assume that the signs are $+1$ and -1. Then the linear system $D^{(k)}\ u_{S_k} = 0$ becomes $T_k \hat{u}_{S_k} = \pm[+1\ 0 \dots\ 0\ -1]^\mathsf{T}$, where T_k is a Toeplitz matrix of suitable dimension and $\hat{u}_{S_k}$ is the vector u_{S_k} without the first and the last entry which we fixed to $+1$ and -1, respectively. The existence of a suitable solution to the linear system $T_k \hat{u}_{S_k}$, which is such that $\|T_k \hat{u}_{S_k}\|_\infty \leq 1$ follows from Lemma 27(ii).

Let us prove the reverse implication, i.e., $z \in \mathcal{S}^\star \Rightarrow z \in \mathcal{SC}$. Without loss of generality, we consider a single segment S_k and we re-label the corresponding entries from 1 to n_k. Let us assume that $z \in \mathcal{S}^\star$, which means that there exists u_{S_k} such that $D^{(k)} u_{S_k} = 0$, $u_{\mathcal{I}_k} = \text{sign}(Dz)_{\mathcal{I}_k}$, and $\|u_{S_k}\|_\infty \leq 1$. Any solution of $D^{(k)} u_{S_k} = 0$ is in the null space of $D^{(k)}$, which is spanned by the vectors v_1 and v_2 defined in Lemma 28. Therefore, we write u_{S_k} as $u_{S_k} = \alpha v_1 + \beta v_2$ with $\alpha, \beta \in \mathbb{R}$. Assume that there are indices $i, j \in S_k$, such that $\text{sign}(Dz)_i = +1$ and $\text{sign}(Dz)_j = -1$, therefore it must hold:

$$[\alpha v_1 + \beta v_2]_i = +1 \quad \text{and} \quad [\alpha v_1 + \beta v_2]_j = -1 \tag{59}$$

which, recalling the definitions of v_1 and v_2 in Lemma 28, becomes: $\alpha + \beta i = +1$ and $\alpha + \beta j = -1$. It follows that:

$$\beta = \frac{2}{j - i} \tag{60}$$

Now, since $z \in \mathcal{S}^\star$ it must also hold $\|u_{S_k}\|_\infty \leq 1$ which can be written as:

$$-\mathbf{1} \leq \quad \alpha v_1 + \beta v_2 \quad \leq \mathbf{1} \Leftrightarrow \tag{61}$$

$$1 \leq \quad \alpha + \beta h \quad \leq 1, \forall h = 1, \dots, n_k \tag{62}$$

Combining the inequalities for $h = 1$ and $h = n_k$ we get:

$$\beta(n_k - 1) \leq 2 \tag{63}$$

Substituting (60) into (63), we get:

$$\frac{n_k - 1}{j - i} \leq 1 \tag{64}$$

which is satisfied if and only if $i = 1$ and $j = n_k$. Hence, we proved that sign changes of the curvature of $z \in \mathcal{S}^\star$ can only happen at the boundary of each segment, which agrees with our definition of $\mathcal{SC}$ (Definition 11), proving the claim. □

G. Proof of Proposition 14

Let us start by proving that if we sample the boundary of $z^\diamond$ and the sample set includes a twin sample in each linear segment of $z^\diamond$, then, $z^\diamond$ is in the set of minimizers of (L1_D). The claim is a direct consequence of Theorem 13. By construction, if we have a twin sample per segment, only two cases are possible: either both samples fall inside the linear segment (i.e., none of them corresponds to a corner), or one of the samples corresponds to a corner. In both cases it is easy to see that $z^\diamond$ is sign consistent with respect to this choice of samples, which implies that $z^\diamond \in \mathcal{S}^\star$ using Theorem 13.

The second claim states that any optimal solution $z^\star$ lies between the naive estimate $\tilde{z}$ and the true profile $z^\diamond$. Before proving this claim we note that the presence of twin samples makes the objective of (L1_D) separable. To see this we note that, in a noiseless case, if we sample a point i, then its value is fixed by the corresponding linear constraint and z_i is no longer a variable. Therefore, by sampling, we are essentially fixing pairs of consecutive entries in z. Using this property, we see that:

$$\|Dz\|_1 = \|D^{(1)} z_{S_1}\|_1 + \|D^{(2)} z_{S_2}\|_1 + \dots + \|D^{(K)} z_{S_K}\|_1 \tag{65}$$

where $D^{(k)}$ is a 2nd-order difference matrix of suitable dimensions, and z_{S_k} is the subvector of z including the entries corresponding to consecutive twin samples (say z_{i-1}, z_i and z_j, z_{j+1}, which are fixed to know values) and all

the entries between those (i.e., $z_{i+1}, \dots, z_{j-1}$); we used the symbol K in (65) to denote the number of regions between consecutive twin samples. From the separability of the objective, it follows that we can study each region (between twin samples) independently (the optimization splits in K independent optimizations). We now prove the second claim: for any optimal solution $z^\star \in \mathcal{S}^\star$ and any index $i \in \{1, \dots, n\}$, it holds that if $z_i^\diamond \leq \tilde{z}_i$ then $z_i^\diamond \leq z_i^\star \leq \tilde{z}_i$, and if $\tilde{z}_i \leq z_i^\diamond$ then $\tilde{z}_i \leq z_i^\star \leq z_i^\diamond$. As mentioned before, this means that any optimal solution is "between" the naive solution $\tilde{z}$ (obtained by connecting the dots, see the blue dashed line in Fig. 5(b)) and the true solution $z^\diamond$ (black solid line in Fig. 5(b)). We show that the claim must hold true in all regions $S_1, \dots, S_K$. First, let us get rid of the "degenerate" regions: these are the ones in which $\tilde{z}_i = z_i^\diamond$ for all $i \in S_k$. This happens when we sample a corner and there are 3 collinear samples as in Fig. 5(a). In this case for any index $i \in S_k$, we prove $\tilde{z}_i = z_i^\star = z_i^\diamond$, i.e., all optimal solutions must reduce to a straight line between the collinear samples. We prove this with the visual support of Fig. 5(a). Our goal is to show that any $\hat{z}$ that deviates from linearity is not sign consistent (SC), hence cannot be optimal. If a sample $k \in \{i+1, \dots, j-1\}$ has a curvature different from zero, then, to be sign consistent it cannot change curvature. This case is shown with the label (↓) in Fig. 5(a). Clearly, if $\hat{z}$ cannot change curvature, after deviating from the straight line, it cannot reach the sample j, leading to contradiction. A similar, contradiction occur when the curvature is different from zero at i: a positive curvature at i (case (↑↓) in figure), must be compensated by a negative curvature before j for the curve to intersect j: this again violates sign consistency; analogous argument holds for a negative curvature at i (case (↓↑) in figure). Therefore we proved that in these straight segments it holds $\tilde{z}_i = z_i^\star = z_i^\diamond$. If only remains to discuss the case in which $\tilde{z}_i \neq z_i^\diamond$ which occurs whenever consecutive double samples do not include corners. This situation is pictured in Fig. 5(b); in this case $\tilde{z}_i \leq z_i^\diamond$. We only prove that whenever $\tilde{z}_i \leq z_i^\diamond$ then $\tilde{z}_i \leq z_i^\star \leq z_i^\diamond$; the proof for the case $z_i^\diamond \leq \tilde{z}_i$ is practically identical (the corner points downwards). In Fig. 5(b) we show two cases (dashed lines): in case (↑↓) we show a signal $\hat{z}$ above $z^\diamond$; in (↓↑) we show a signal below $\tilde{z}$. One can easily realize than any signal as in case (↑↓) has a positive curvature at sample i and a negative curvature at the top corner: this violates sign consistency hence $\hat{z}$ is not in the solution set. A similar argument holds in the case (↓↑), which concludes the proof the $z^\star$ must be "between" $\tilde{z}$ and $z^\diamond$.

Extra result: error bounds for $z^\star$. We prove the following error bound which relates any solution $z^\star$ of the ℓ_1-minimization problem (L1$_D$) to the true profile $z^\diamond$:

$$\|z^\diamond - z^\star\|_\infty \leq \max_{i \in \mathcal{M}} d_i \cos(\theta_i) \tag{66}$$

where d_i is the distance between the sample i and the nearest corner in $z^\diamond$, while θ_i is the angle that the line connecting i with the nearest corner forms with the vertical, see Fig. 5(b).

To prove (66), we note that $z^\star$ can deviate from $z^\diamond$ only in the "corner cases" as the one in Fig. 5(b) (proven above in this section). Moreover, we note that the maximum error $|z_i^\diamond - \tilde{z}_i|$ is attained at the corner of $z^\diamond$ and is denoted with $\hat{d}$ in the figure. From basic trigonometry we conclude:

$$\hat{d} \leq \max_{k \in \{i,j\}} d_k \cos\theta_k \tag{67}$$

(in Fig. 5 this becomes: $\hat{d} \leq d_i \cos(\theta_i)$), where d_i is the distance between sample i and the nearest corner, while θ_i is the angle that the line connecting sample i with the nearest corner forms with the vertical. The bound (66) follows by extending this inequality to all regions in $z^\diamond$. □

H. Proof of Theorem 17

In this section we prove that if a 3D profile Z is feasible for problem (L1$_\Delta$) and it is 3D sign consistent, then it is also a minimizer of (L1$_\Delta$). The proof is similar to the 2D case in Proposition 14, and relies on Corollary 10, which we restate as follows: given a profile $Z \in \mathbb{R}^{r \times c}$ which is feasible for problem (L1$_\Delta$), Z is in the optimal set of (L1$_\Delta$) if there exists a vector $u \in \mathbb{R}^{2(n-r-c)}$, with $n = r \times c$, such that

$$(\Delta^\mathsf{T})_{\overline{\mathcal{M}}}\, u = \mathbf{0} \;\text{ and }\; u_\mathcal{I} = \text{sign}(\Delta z)_\mathcal{I} \;\text{ and }\; \|u_\mathcal{J}\|_\infty \leq 1, \tag{68}$$

where $\mathcal{M}$ denotes a grid sample set. Let $\overline{\mathcal{M}}_i$ be patches defined in Definition 15. We show that 3D sign consistency of Z with respect to grid samples $\mathcal{M}$ implies (68). We start by noting that when using grid samples with K patches (see Fig. 7), problem (L1$_\Delta$) separates into K independent optimization subproblems (similarly to the 2D case of eq. (65)). Therefore, without loss of generality in the following we focus on a single patch and we assume that the grid samples include the boundaries of the patch (first and last two rows and columns). With slight abuse of notation we denote this patch with Z and we use $z = \text{vec}(Z)$.

Before proving the claim we need some insight on the structure of the matrix $(\Delta^\mathsf{T})_{\overline{\mathcal{M}}}$. This matrix is obtained by deleting rows of Δ^T indexed by $\mathcal{M}$. Since we are assuming to sample the boundaries of the patch, the resulting $(\Delta^\mathsf{T})_{\overline{\mathcal{M}}}$ has the structure described in Fig. 27. Now we remain to show that when Z is 3D sign consistent, we can find a vector u that satisfies the three conditions in (68). Towards this goal, we recall that Δ is obtained by stacking two submatrices that compute the vertical and horizontal differences as in (11). Therefore, we split $\text{sign}(\Delta z)_\mathcal{I}$ accordingly as:

$$\text{sign}(\Delta z)_\mathcal{I} = \begin{bmatrix} \text{sign}(\Delta z)_{\mathcal{I}_V} \\ \text{sign}(\Delta z)_{\mathcal{I}_H} \end{bmatrix} \tag{69}$$

where $\mathcal{I}_V$ includes indices in the support set $\mathcal{I}$ corresponding to nonzero vertical differences, while $\mathcal{I}_H$ includes indices in $\mathcal{I}$ corresponding to nonzero horizontal differences.

Now, we note that the row partition of Δ induces a column partition of $(\Delta^\mathsf{T})_{\overline{\mathcal{M}}}$. We call the corresponding submatrices $(\Delta^\mathsf{T})_{\overline{\mathcal{M}},V}$ and $(\Delta^\mathsf{T})_{\overline{\mathcal{M}},H}$, as shown in Fig. 27. This also partitions the vector u into two subvectors u_V and u_H. Therefore, we rewrite the condition $(\Delta^\mathsf{T})_{\overline{\mathcal{M}}} u = \mathbf{0}$ as:

$$(\Delta^\mathsf{T})_{\overline{\mathcal{M}},V}\, u_V + (\Delta^\mathsf{T})_{\overline{\mathcal{M}},H}\, u_H = \mathbf{0} \tag{70}$$

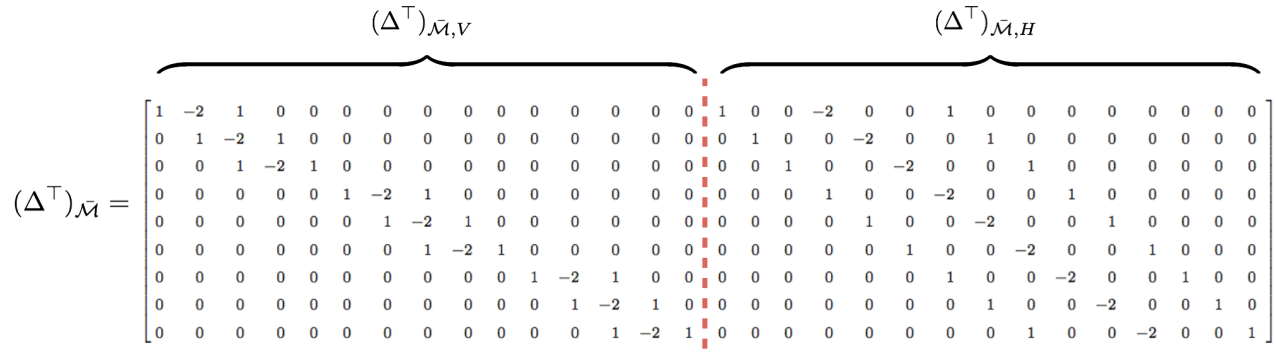

Fig. 27. Example of matrix $(\Delta^\mathsf{T})_{\overline{\mathcal{M}}}$ and corresponding column partition $(\Delta^\mathsf{T})_{\overline{\mathcal{M}},V}$, $(\Delta^\mathsf{T})_{\overline{\mathcal{M}},H}$.

Since Z is 3D sign consistent, then $\text{sign}(\Delta z)_{\mathcal{I}_V}$ is either $+\mathbf{1}$ or $-\mathbf{1}$ and $\text{sign}(\Delta z)_{\mathcal{I}_H}$ is either $+\mathbf{1}$ or $-\mathbf{1}$. Assume without loss of generality that $\text{sign}(\Delta z)_{\mathcal{I}_V} = +\mathbf{1}$ and $\text{sign}(\Delta z)_{\mathcal{I}_H} = -\mathbf{1}$. Now if we choose $u_V = +\mathbf{1}$ and $u_H = -\mathbf{1}$, it holds that $u_{\mathcal{I}_V} = \text{sign}(\Delta z)_{\mathcal{I}_V}$, $u_{\mathcal{I}_H} = \text{sign}(\Delta z)_{\mathcal{I}_H}$, and $\|u_{\mathcal{J}}\|_\infty \leq \|u\|_\infty \leq 1$, hence the last two conditions in (68) are satisfied. Moreover, since each row of $(\Delta^\mathsf{T})_{\overline{\mathcal{M}},V}$ and $(\Delta^\mathsf{T})_{\overline{\mathcal{M}},H}$ includes only three nonzero entries with values $+1, -2, +1$, it follows that:

$$\begin{aligned}(\Delta^\mathsf{T})_{\overline{\mathcal{M}},V}\, u_V &= (\Delta^\mathsf{T})_{\overline{\mathcal{M}},V}\, \mathbf{1} = \mathbf{0}\\ (\Delta^\mathsf{T})_{\overline{\mathcal{M}},H}\, u_H &= -(\Delta^\mathsf{T})_{\overline{\mathcal{M}},H}\, \mathbf{1} = \mathbf{0}\end{aligned} \tag{71}$$

which implies (70), concluding the proof. □

I. Proof of Proposition 18

By the assumption of Proposition 18, $Z^\diamond$ is the ground truth generating noiseless measurements (4) and $Z^\diamond$ is *3D sign consistent* with respect to $\mathcal{M}$. Then each row of $Z^\diamond$, namely $Z_i^\diamond$, is a sign-consistent 2D depth profile and, given the samples, we can build a *row-wise envelop* for $Z_i^\diamond$ as prescribed in Theorem 13 (pictorial explanation in Fig. 6(b)). Repeating this procedure for all rows i and calling $\bar{Z}$ and $\underline{Z}$ the point-wise upper bound an lower bound for the envelop of each row, then $Z^\diamond$ is within the *row-wise envelop*, i.e., $\underline{Z}_{i,j} \leq Z^\diamond_{i,j} \leq \bar{Z}_{i,j}$ for all $i = 1, \dots, r$ and $j = 1, \dots, c$. Therefore, given an optimal solution $Z^\star$ of (L1_Δ) it holds that

$$|Z^\diamond_{i,j} - Z^\star_{i,j}| \leq \max(|\underline{Z}_{i,j} - {Z^\star}_{i,j}|, |\bar{Z}_{i,j} - {Z^\star}_{i,j}|).$$

J. Proof of Proposition 19

This section proves necessary and sufficient conditions for an estimate $z^\star$ to be in the set $\mathcal{S}^\star$ of optimal solutions of problem (L1_D^ε). We follow similar arguments as in the previous proof. Rewrite (L1_D^ε) as:

$$\min_z \|Dz\|_1 + \chi_{\{z:\|Az-y\|_\infty \leq \varepsilon\}} =: \min_z f(z) \tag{72}$$

As discussed earlier in Section E, we will establish the conditions so that the zero vector belongs to the set of subgradients of f at $z^\star$, i.e., $\mathbf{0} \in \partial f(z^\star)$. The subdifferential of $\|Dz\|_1$ was given in (46). The subdifferential of the indicator function in (72) is given similarly to (47)

$$\begin{aligned}&\partial(\chi_{\{z:\|Az-y\|_\infty \leq \varepsilon\}})(z^\star) = \\ &\quad \{g \in \mathbb{R}^n : g^\mathsf{T} z^\star \geq g^\mathsf{T} r,\ \forall r \text{ s.t. } \|Ar - y\|_\infty \leq \varepsilon\}\end{aligned} \tag{73}$$

Recall that the matrix A restricts a vector r to its entries in the sample set $\mathcal{M}$, therefore we have

$$\{r :\ \|Ar - y\|_\infty \leq \varepsilon\} = \{r :\ \|r_\mathcal{M} - y\|_\infty \leq \varepsilon\}. \tag{74}$$

Obviously $\|r_\mathcal{M} - y\|_\infty \leq \varepsilon$ implies for all $i \in \mathcal{M}$ that

$$r_i \leq y_i + \varepsilon \quad \text{or} \quad r_i \geq y_i - \varepsilon. \tag{75}$$

We decompose

$$g^\mathsf{T} z^\star \geq g^\mathsf{T} r = g_\mathcal{M}^\mathsf{T} r_\mathcal{M} + g_{\overline{\mathcal{M}}}^\mathsf{T} r_{\overline{\mathcal{M}}}. \tag{76}$$

A vector g that satisfies (76) for all r obeying (74) must satisfy that $g_{\overline{\mathcal{M}}}^\mathsf{T} = \mathbf{0}$ since $r_{\overline{\mathcal{M}}}$ is a free variable. With this at hand, we proceed as

$$\begin{aligned}g^\mathsf{T} z^\star &\geq \max_{\|r_\mathcal{M} - y\|_\infty \leq \varepsilon} g_\mathcal{M}^\mathsf{T} r_\mathcal{M} = \max_{\|r_\mathcal{M} - y\|_\infty \leq \varepsilon} \sum_{i \in \mathcal{M}} g_i r_i \\ &= \max_{\|r_\mathcal{M} - y\|_\infty \leq \varepsilon} \left(\sum_{g_i \geq 0,\ i \in \mathcal{M}} |g_i| r_i - \sum_{g_i < 0,\ i \in \mathcal{M}} |g_i| r_i \right) \\ &= \sum_{g_i \geq 0,\ i \in \mathcal{M}} |g_i|(y_i + \varepsilon) - \sum_{g_i < 0,\ i \in \mathcal{M}} |g_i|(y_i - \varepsilon) \\ &= \sum_{g_i \geq 0,\ i \in \mathcal{M}} |g_i| y_i - \sum_{g_i < 0,\ i \in \mathcal{M}} |g_i| y_i + \varepsilon \sum_{i \in \mathcal{M}} |g_i| \\ &= g_\mathcal{M}^\mathsf{T} y + \varepsilon \|g_\mathcal{M}\|_1.\end{aligned}$$

Rearranging terms yields

$$\|g_\mathcal{M}\|_1 \leq \frac{g_\mathcal{M}^\mathsf{T}(z^\star_\mathcal{M} - y)}{\varepsilon}. \tag{77}$$

Then we can formulate the solution set as

$$\begin{aligned}\mathcal{S}^\star = \{z^\star :\quad &\exists u \in \mathbb{R}^{n-2}, g \in \mathbb{R}^n, \text{ such that} \\ &D^\mathsf{T} u + g = \mathbf{0}, \\ &u_\mathcal{I} = \text{sign}(Dz^\star)_\mathcal{I},\ \ \|u_\mathcal{J}\|_\infty \leq 1, \\ &\|g_\mathcal{M}\|_1 \leq \frac{g_\mathcal{M}^\mathsf{T}(z^\star_\mathcal{M} - y)}{\varepsilon}, \\ &g_{\overline{\mathcal{M}}}^\mathsf{T} = \mathbf{0}\}.\end{aligned} \tag{78}$$

The conditions $D^{\mathsf{T}}u + g = \mathbf{0}$ and $g_{\overline{\mathcal{M}}}^{\mathsf{T}} = \mathbf{0}$ can be cast as $(D^{\mathsf{T}})_{\overline{\mathcal{M}}}\, u = \mathbf{0}$ and $g_{\mathcal{M}} = -(D^{\mathsf{T}})_{\mathcal{M}}\, u$. Then inequality in (77) must be an equality since $|z^\star_{\mathcal{M}} - y| \leq \varepsilon$ and $\|g_{\mathcal{M}}\|_1 = g_{\mathcal{M}}^{\mathsf{T}}\mathrm{sign}(g_{\mathcal{M}})$ by definition of the ℓ_1-norm and moreover if $|z_i^\star - y_i| = \varepsilon$ for some $i \in \mathcal{M}$ it must hold that

$$\mathrm{sign}((D^{\mathsf{T}})_i\, u) = (z_i^\star - y_i)/\varepsilon. \tag{79}$$

Here we used the assumption that $\varepsilon > 0$. Denote the set of such $i \in \mathcal{M}$ as the active set $\mathcal{A}$ as in (16). For $i \in \mathcal{M} \setminus \mathcal{A}$, it holds that $(D^{\mathsf{T}})_i\ u = 0$. Splitting $\mathcal{A}$ into two subsets as in (17) and using (79) yields the conditions (19) of Proposition 19.

K. Proof of Proposition 23

From Theorem 21 we learn that any optimal solution $z^\star$ must be 2D sign consistent, implyting that it must also lie within the 2D envelop. By the assumption of Proposition 23 that we sample the boundary of $z^\diamond$ and the sample set includes a twin sample in each linear segment in $z^\diamond$, it follows that the ground truth signal $z^\diamond$ is also 2D sign consistent. Therefore, the gap between $z^\star$ and $z^\diamond$ is no more than $\bar{Z}_{i,j} - \underline{Z}_{i,j}$.

With simple geometric argument (see Fig. 9), we can readily derive that

$$\bar{Z}_{i,j} - \underline{Z}_{i,j} = \leq d_i \cos(\theta_i) + d_j \cos(\theta_j) + 2\varepsilon.$$

L. Proof of Corollary 25

In the proof of Proposition 14 we have seen that the objective of (L1$_D$) is separable as in (65), where S_k includes all indices falling in the region between two consecutive twin samples (e.g., if the twin samples are $(i-1, i)$ and $(j, j+1)$, then $S_k = \{i-1, i, i+1, \ldots, j-1, j, j+1\}$). This allows us to study the performance of Algorithm 1 independently for each region S_k. As a result of Proposition 14, line 2 in Algorithm 1 produces a solution $z^\star$ between $\tilde{z}$ and $z^\diamond$. Assume without loss of generality that $z^\star$ is concave in the region S_k, e.g., the leftmost corner of [**?**, Fig. 2(b)], which yields $s_k = -1$ in line 5. Then $\tilde{z}_i \leq z_i^\star \leq z_i^\diamond$ for all $i \in S_k$ which implies

$$\sum_{i \in S_k} z_i^\star \leq \sum_{i \in S_k} z_i^\diamond \tag{80}$$

Now we note that maximizing $\sum_{i \in S_k} z_i$ is the same as minimizing $-\sum_{i \in S_k} z_i$. Therefore, when $s_k = -1$, line 6 maximizes the objective $\sum_{i \in S_k} z_i$ subject to $z \in \mathcal{S}^\star$ where $\mathcal{S}^\star$ is the optimal set of (L1$_D$). Since (80) holds for all $z \in \mathcal{S}^\star$, and since $z^\diamond \in \mathcal{S}^\star$, Algorithm 1 returns $z^\diamond$. □

M. Computations for NESTA solver

In this section we provide some technical derivations for the Algorithm 2. The function f_μ in (22) is shown to be differentiable by Nesterov [48] and has the gradient

$$\nabla f_\mu(z) = D^{\mathsf{T}} u^\star(z) \tag{81}$$

where $u^\star(z)$ is the optimal solution to (22) and is given by

$$u^\star(z) = \begin{cases} \mu^{-1}(Dz)_i, & \text{if } |(Dz)_i| < \mu \\ \mathrm{sign}(Dz)_i, & \text{otherwise.} \end{cases} \tag{82}$$

The gradient $\nabla f_\mu(z)$ is said to be Lipschitz with Lipschitz constant L_μ if it obeys

$$\|\nabla f_\mu(x) - \nabla f_\mu(z)\|_2 \leq L_\mu \|x - z\|_2.$$

This constant is used in Algorithm 2 and is shown to be $L_\mu = \frac{\|D\|^2}{\mu}$ [4, eq (3.4)] where $\|\cdot\|$ denotes the spectral norm.

Next we move onto solving lines 8-10 in Algorithm 2. We start with $x^{(k)}$. A bit of algebra yields

$$\begin{aligned} x^{(k)} =& \mathrm{argmin}_{z \in \mathcal{Q}}\ \frac{L_\mu}{2}\|z - z^{(k)}\|_2^2 + \langle c, z\rangle \\ =& \mathrm{argmin}_{z \in \mathcal{Q}} \|z - (z^{(k)} - L_\mu^{-1} c)\|_2^2 \end{aligned} \tag{83}$$

where $c = \nabla f_\mu(z^{(k)})$ and $\mathcal{Q} = \{z : \|Az - y\|_\infty \leq \varepsilon\}$. For simplicity we denote $\hat{z}^{(k)} = z^{(k)} - L_\mu^{-1} c$ and recall that the sampling matrix A restricts a vector to its entries in a sampling set $\mathcal{M}$. Therefore we can write $\mathcal{Q} = \{z : \|z_{\mathcal{M}} - y\|_\infty \leq \varepsilon\}$. This suggests that the optimal solution $x^{(k)}$ must agree with $z^{(k)}$ on $\overline{\mathcal{M}}$, i.e., $x^{(k)}_{\overline{\mathcal{M}}} = \hat{z}^{(k)}_{\overline{\mathcal{M}}}$. It remains to solve

$$x^{(k)}_{\mathcal{M}} = \mathrm{argmin}_z \|z - \hat{z}^{(k)}_{\mathcal{M}}\|_2^2 \quad \text{subject to } \|z - y\|_\infty \leq \varepsilon. \tag{84}$$

The problem above is nothing but an Euclidean projection of the vector $\hat{z}^{(k)}_{\mathcal{M}}$ onto the ℓ_∞-ball with radius ε centered at y. Assuming that $\mathcal{M} = \{1, 2, \ldots, |\mathcal{M}|\}$ without loss of generality, this projection can be explicitly calculated as follows:

$$x_i^{(k)} = y_i + \begin{cases} \varepsilon, & \text{if } r_i > \varepsilon \\ -\varepsilon, & \text{if } r_i < -\varepsilon \\ r_i, & \text{if } |r_i| \leq \varepsilon \end{cases}, \quad i \in \mathcal{M}$$

where $r = \hat{z}^{(k)}_{\mathcal{M}} - y$. We obtained a closed form solution for $x^{(k)}$. The same calculation holds for $y^{(k)}$ as well by replacing $c = \sum_{i=0}^{k} \alpha_i \nabla f_\mu(z^{(i)})$ in (83).

N. Extra visualizations: Gazebo depth images

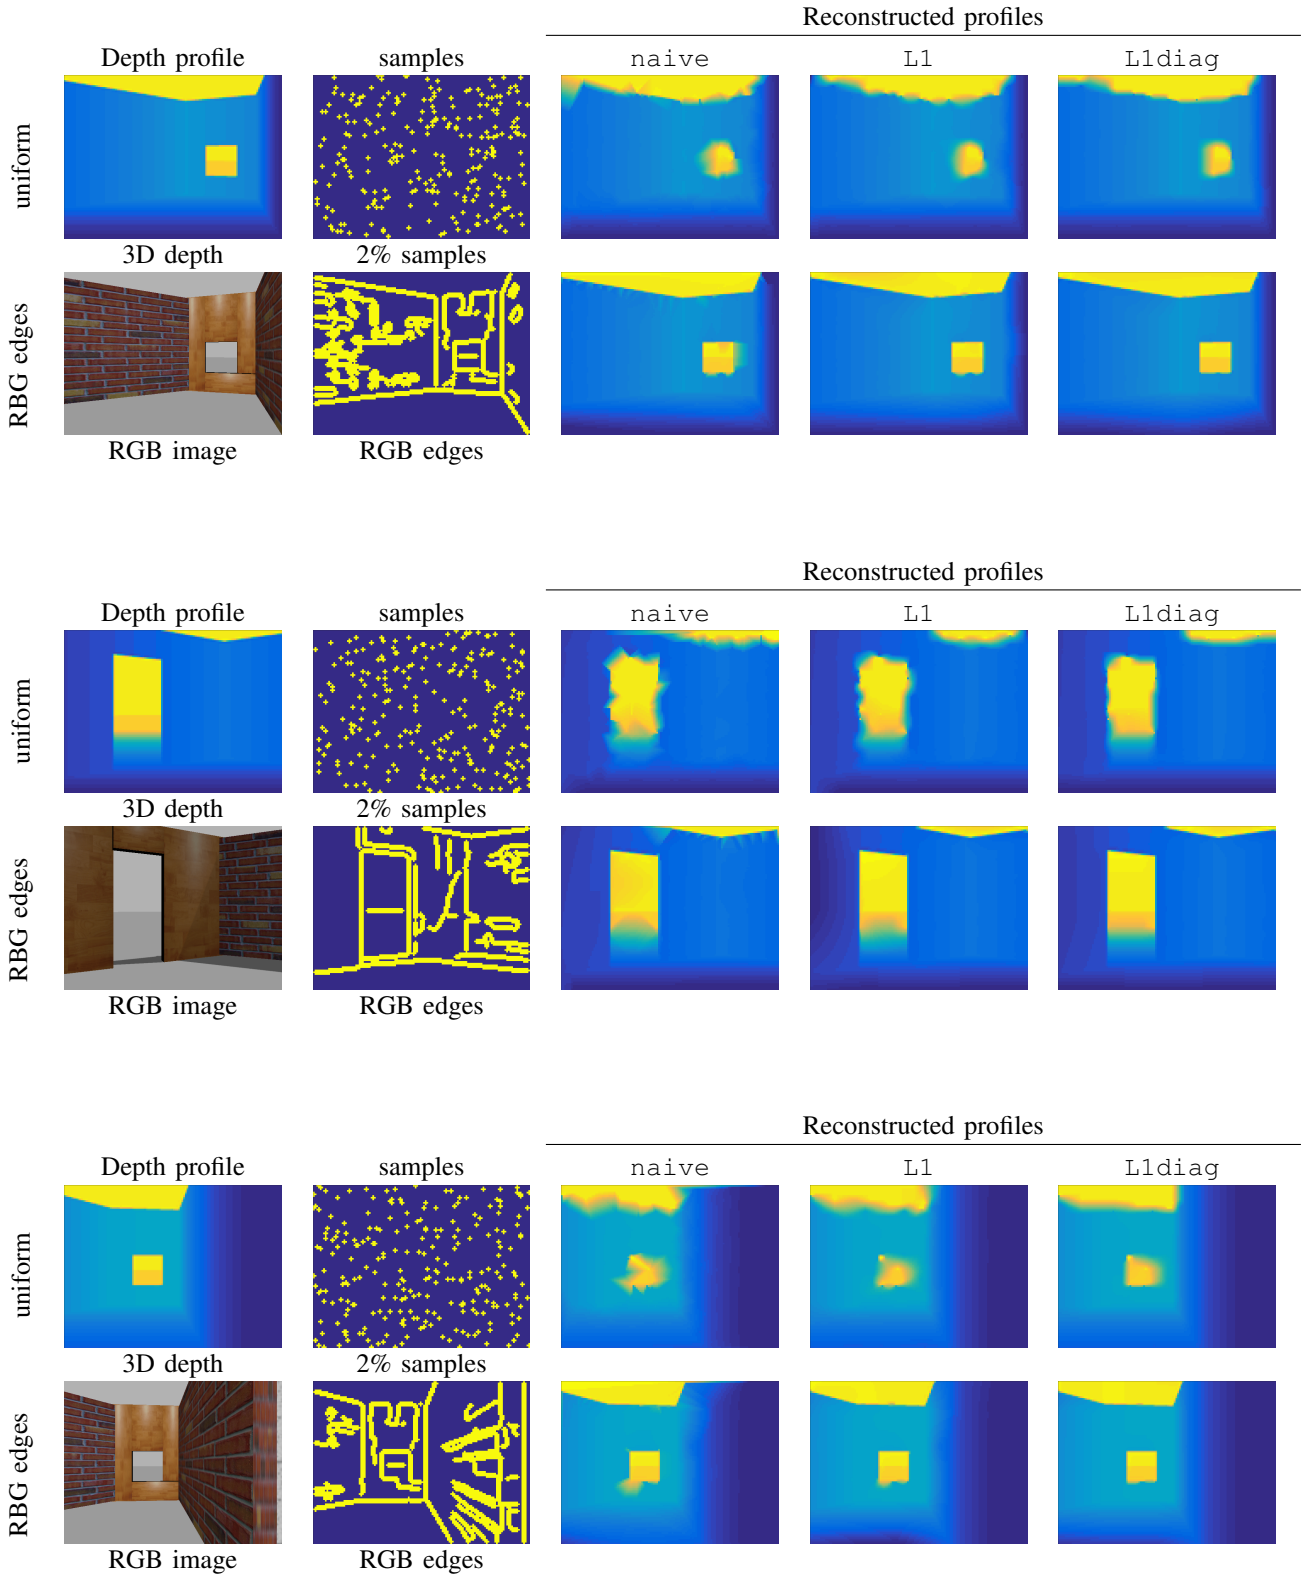

Fig. 28. Gazebo: 3 examples of reconstructed depth profiles using the proposed approaches (`L1`, `L1diag`) and a naive linear interpolation (`naive`). For each example we show the reconstruction from 2% uniformly drawn depth measurements. We also show the reconstruction for the case in which we can access the depth corresponding to (appearance) edges in the RGB images.

O. Extra visualizations: ZED depth images

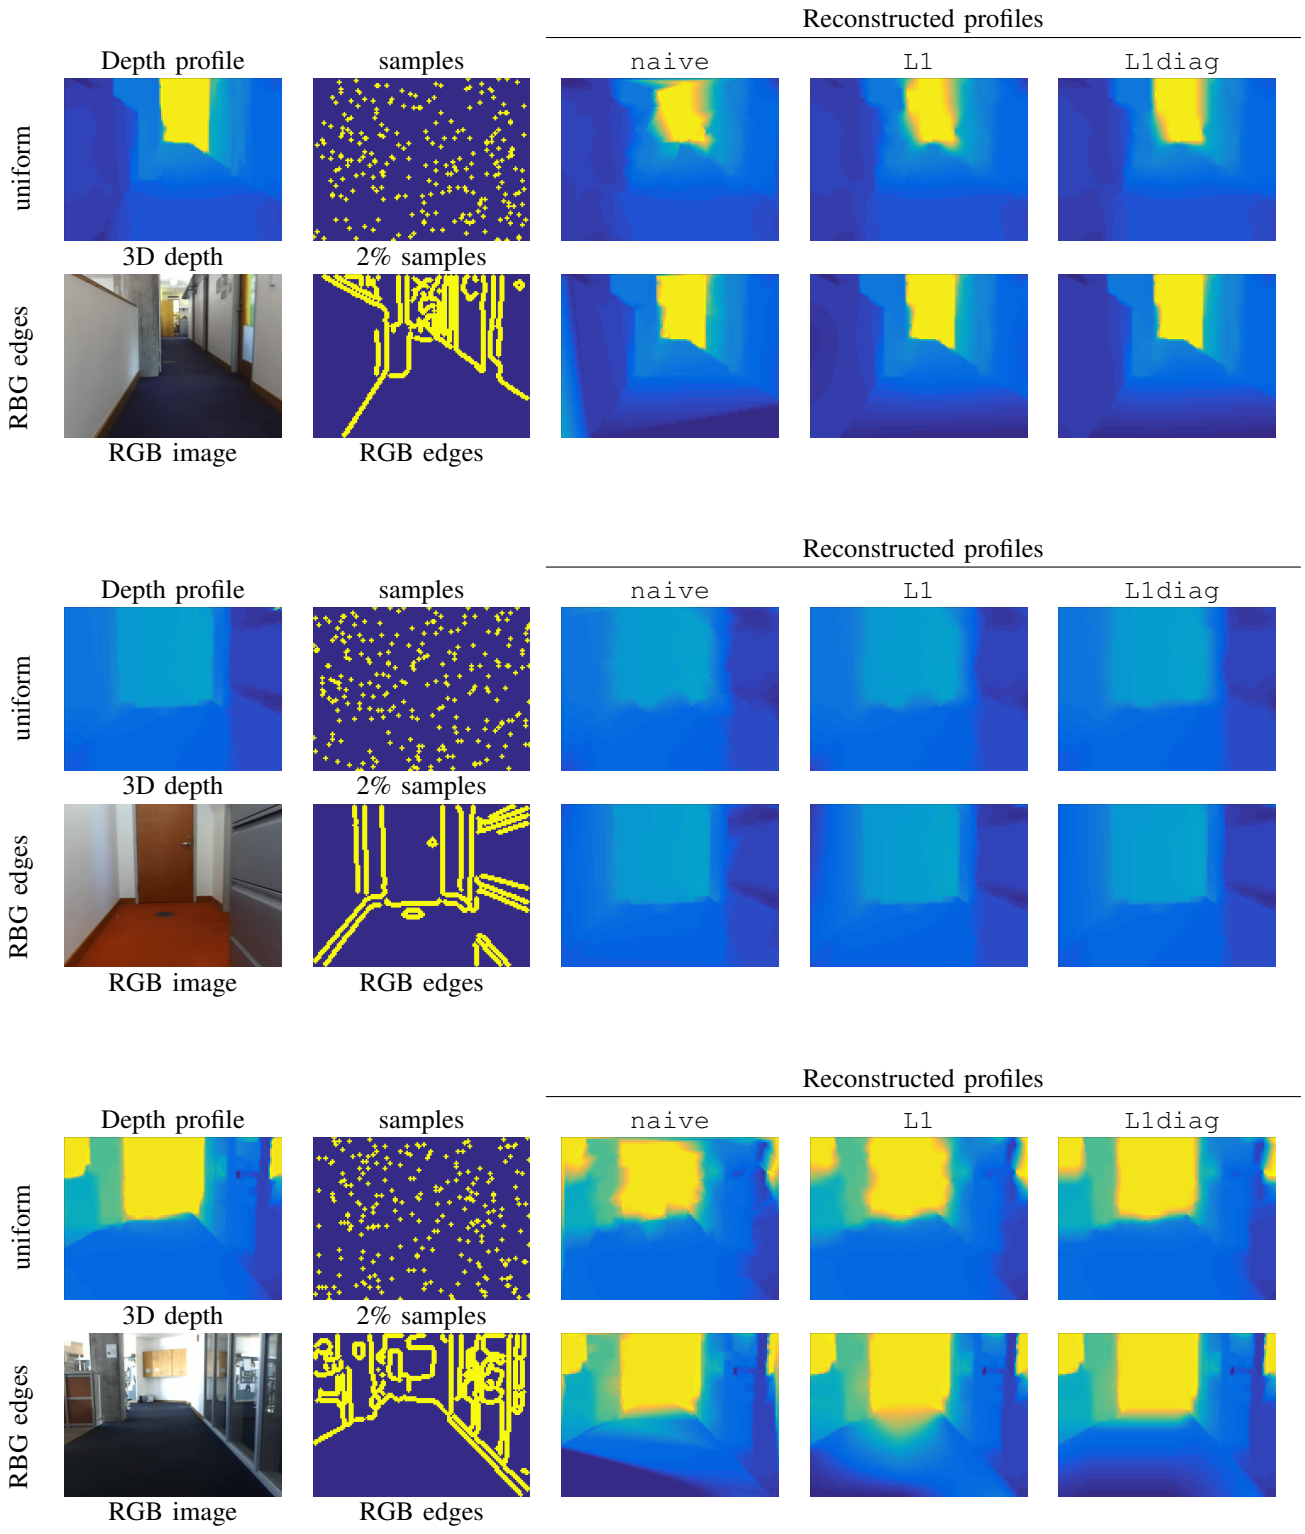

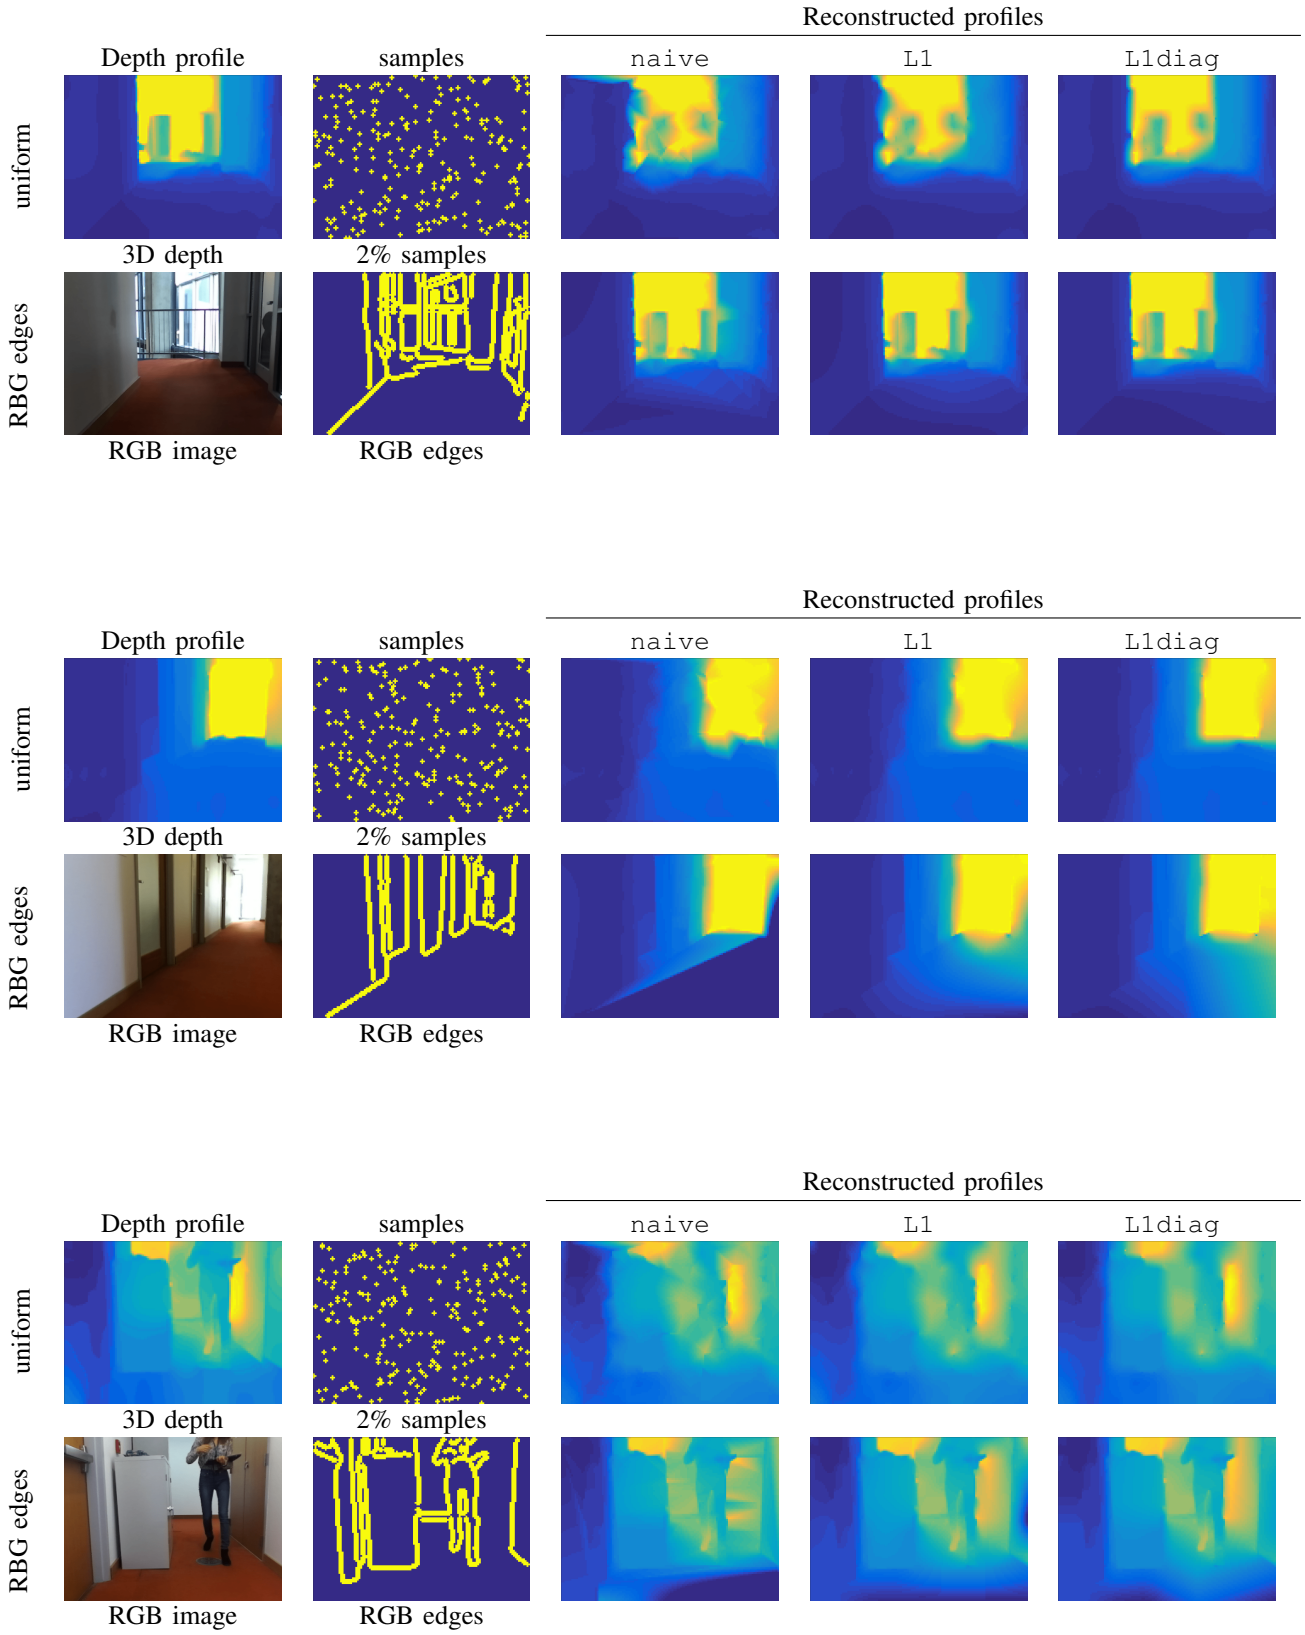

Fig. 29. ZED: 6 examples of reconstructed depth profiles using the proposed approaches (`L1`, `L1diag`) and a naive linear interpolation (`naive`). For each example we show the reconstruction from 2% uniformly drawn depth measurements. We also show the reconstruction for the case in which we can access the depth corresponding to (appearance) edges in the RGB images.